



04/26/00

00000000000000000000000000000000

04-28-00

A

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Inventorship..... Puryear
Applicant..... Microsoft Corporation
Attorney's Docket No. MS1-540US
Title: Extensible Kernel-Mode Audio Processing Architecture

JC542 U.S. PTO
09/559901
04/26/00TRANSMITTAL LETTER AND CERTIFICATE OF MAILING

To: Commissioner of Patents and Trademarks,
Washington, D.C. 20231

From: Allan T. Sponseller (Tel. 509-324-9256; Fax 509-323-8979)
Lee & Hayes, PLLC
421 W. Riverside Avenue, Suite 500
Spokane, WA 99201

The following enumerated items accompany this transmittal letter and are being submitted for the matter identified in the above caption.

1. Specification—title page, plus 60 pages, including 69 claims and Abstract
2. Transmittal letter including Certificate of Express Mailing
3. 10 Sheets Formal Drawings (Figs. 1-12)
4. Return Post Card

Large Entity Status [x]

Small Entity Status []

Date: April 26, 2000

By:

Allan T. Sponseller
Reg. No. 38,318CERTIFICATE OF MAILING

I hereby certify that the items listed above as enclosed are being deposited with the U.S. Postal Service as either first class mail, or Express Mail if the blank for Express Mail No. is completed below, in an envelope addressed to The Commissioner of Patents and Trademarks, Washington, D.C. 20231, on the below-indicated date. Any Express Mail No. has also been marked on the listed items.

Express Mail No. (if applicable) EL580804334Date: 4.26.2000

By:

Lori A. Vierra

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

Extensible Kernel-Mode Audio Processing Architecture

Inventor(s):

Martin G. Puryear

ATTORNEY'S DOCKET NO. MS1-540US



04/26/00

1

2 **RELATED APPLICATIONS**3 This application claims the benefit of U.S. Provisional Application No.
4 60/197,100, filed April 12, 2000, entitled "Extensible Kernel-Mode Audio
5 Processing Architecture" to Martin G. Puryear.

6

7 **TECHNICAL FIELD**8 This invention relates to audio processing systems. More particularly, the
9 invention relates to an extensible kernel-mode audio processing architecture.

10

11 **BACKGROUND OF THE INVENTION**12 Musical performances have become a key component of electronic and
13 multimedia products such as stand-alone video game devices, computer-based
14 video games, computer-based slide show presentations, computer animation, and
15 other similar products and applications. As a result, music generating devices and
16 music playback devices are now tightly integrated into electronic and multimedia
17 components.18 Musical accompaniment for multimedia products can be provided in the
19 form of digitized audio streams. While this format allows recording and accurate
20 reproduction of non-synthesized sounds, it consumes a substantial amount of
21 memory. As a result, the variety of music that can be provided using this approach
22 is limited. Another disadvantage of this approach is that the stored music cannot
23 be easily varied. For example, it is generally not possible to change a particular
24 musical part, such as a bass part, without re-recording the entire musical stream.

25

Because of these disadvantages, it has become quite common to generate music based on a variety of data other than pre-recorded digital streams. For example, a particular musical piece might be represented as a sequence of discrete notes and other events corresponding generally to actions that might be performed by a keyboardist—such as pressing or releasing a key, pressing or releasing a sustain pedal, activating a pitch bend wheel, changing a volume level, changing a preset, etc. An event such as a note event is represented by some type of data structure that includes information about the note such as pitch, duration, volume, and timing. Music events such as these are typically stored in a sequence that roughly corresponds to the order in which the events occur. Rendering software retrieves each music event and examines it for relevant information such as timing information and information relating the particular device or “instrument” to which the music event applies. The rendering software then sends the music event to the appropriate device at the proper time, where it is rendered. The MIDI (Musical Instrument Digital Interface) standard is an example of a music generation standard or technique of this type, which represents a musical performance as a series of events.

Computing devices, such as many modern computer systems, allow MIDI data to be manipulated and/or rendered. These computing devices are frequently built based on an architecture employing multiple privilege levels, often referred to as user-mode and kernel-mode. Manipulation of the MIDI data is typically performed by one or more applications executing in user-mode, while the input of data from and output of data to hardware is typically managed by an operating system or a driver executing in kernel-mode.

Such a setup requires the MIDI data to be received by the driver or operating system executing in kernel-mode, transferred to the application executing in user-mode, manipulated by the application as needed in user-mode, and then transferred back to the operating system or driver executing in kernel-mode for rendering. Data transfers between kernel-mode and user-mode, however, can take a considerable and unpredictable amount of time. Lengthy delays can result in unacceptable latency, particularly for real-time audio playback, while unpredictability can result in an unacceptable amount of jitter in the audio data, resulting in unacceptable rendering of the audio data.

The invention described below addresses these disadvantages, providing an extensible kernel-mode audio processing architecture.

SUMMARY OF THE INVENTION

An extensible kernel-mode audio processing architecture is described herein.

According to one aspect, an audio processing architecture is implemented using multiple modules that together form a module graph. The module graph is implemented in kernel-mode, reducing latency and jitter when handling audio data by avoiding transfers of the audio data to user-mode applications for processing.

According to another aspect, the audio processing architecture is a MIDI data processing architecture.

According to another aspect, an interface is described for implementation on each of the multiple modules in a module graph. The interface provides a relatively quick and low-overhead interface for kernel-mode modules to communicate audio data to one another. The interface includes a ConnectOutput

1 interface via which the next module in the graph (that is, the module that audio
2 data should be output to) can be identified to the module, and a DisconnectOutput
3 interface via which the previously-set next module can be cleared (e.g., to a
4 default value, such as an allocator module). The interface also includes a
5 PutMessage interface which is called to pass audio packets to the next module in
6 the graph, and a SetState interface which is called to set the state of the module
7 (e.g., run, stop, or a transitional pause or acquire state).

8 According to another aspect, the audio processing architecture is readily
9 extensible. The audio processing architecture is implemented as multiple kernel-
10 mode modules connected together in a module graph by a graph builder. The
11 graph builder can readily change the module graph, adding new modules,
12 removing modules, or altering connections as necessary, all while the graph is
13 running.

14 According to another aspect, the audio processing architecture includes an
15 allocator that allocates memory for data packets that are passed among modules in
16 a kernel-mode module graph. The allocated memory can be on a data packet
17 basis, or alternatively larger buffers may be allocated to accommodate larger
18 portions of audio data.

19

20 **BRIEF DESCRIPTION OF THE DRAWINGS**

21 The present invention is illustrated by way of example and not limitation in
22 the figures of the accompanying drawings. The same numbers are used
23 throughout the figures to reference like components and/or features.

24 Fig. 1 is a block diagram illustrating an exemplary system for manipulating
25 and rendering audio data.

1 Fig. 2 shows a general example of a computer that can be used in
2 accordance with certain embodiments of the invention.

3 Fig. 3 is a block diagram illustrating an exemplary MIDI processing
4 architecture in accordance with certain embodiments of the invention.

5 Fig. 4 is a block diagram illustrating an exemplary transform module graph
6 module in accordance with certain embodiments of the invention.

7 Fig. 5 is a block diagram illustrating an exemplary MIDI message.

8 Fig. 6 is a block diagram illustrating an exemplary MIDI data packet in
9 accordance with certain embodiments of the invention.

10 Fig. 7 is a block diagram illustrating an exemplary buffer for
11 communicating MIDI data between a non-legacy application and a MIDI
12 transform module graph module in accordance with certain embodiments of the
13 invention.

14 Fig. 8 is a block diagram illustrating an exemplary buffer for
15 communicating MIDI data between a legacy application and a MIDI transform
16 module graph module in accordance with certain embodiments of the invention.

17 Fig. 9 is a block diagram illustrating an exemplary MIDI transform module
18 graph such as may be used in accordance with certain embodiments of the
19 invention.

20 Fig. 10 is a block diagram illustrating another exemplary MIDI transform
21 module graph such as may be used in accordance with certain embodiments of the
22 invention.

23 Fig. 11 is a flowchart illustrating an exemplary process for the operation of
24 a module in a MIDI transform module graph in accordance with certain
25 embodiments of the invention.

1 Fig. 12 is a flowchart illustrating an exemplary process for the operation of
2 a graph builder in accordance with certain embodiments of the invention.

3

4 **DETAILED DESCRIPTION**

5 **General Environment**

6 Fig. 1 is a block diagram illustrating an exemplary system for manipulating
7 and rendering audio data. One type of audio data is defined by the MIDI (Musical
8 Instrument Digital Interface) standard, including both accepted versions of the
9 standard and proposed versions for future adoption. Although various
10 embodiments of the invention are discussed herein with reference to the MIDI
11 standard, other audio data standards can alternatively be used. In addition, other
12 types of audio control information can also be passed, such as volume change
13 messages, audio pan change messages (e.g., changing the manner in which the
14 source of sound appears to move from two or more speakers), a coordinate change
15 on a 3D sound buffer, messages for synchronized start of multiple devices, or any
16 other parameter of how the audio is being processed.

17 Audio system 100 includes a computing device 102 and an audio output
18 device 104. Computing device 102 represents any of a wide variety of computing
19 devices, such as conventional desktop computers, gaming devices, Internet
20 appliances, etc. Audio output device 104 is a device that renders audio data,
21 producing audible sounds based on signals received from computing device 102.
22 Audio output device 104 can be separate from computing device 102 (but coupled
23 to device 102 via a wired or wireless connection), or alternatively incorporated
24 into computing device 102. Audio output device 104 can be any of a wide variety

1 of audible sound-producing devices, such as an internal personal computer
2 speaker, one or more external speakers, etc.

3 Computing device 102 receives MIDI data for processing, which can
4 include manipulating the MIDI data, playing (rendering) the MIDI data, storing
5 the MIDI data, transporting the MIDI data to another device via a network, etc.
6 MIDI data can be received from a variety of devices, examples of which are
7 illustrated in Fig. 1. MIDI data can be received from a keyboard 106 or other
8 musical instruments 108 (e.g., drum machine, synthesizer, etc.), another audio
9 device(s) 110 (e.g., amplifier, receiver, etc.), a local (either fixed or removable)
10 storage device 112, a remote (either fixed or removable) storage device 114,
11 another device 116 via a network (such as a local area network or the Internet),
12 etc. Some of these MIDI data sources can generate MIDI data (e.g., keyboard
13 106, audio device 110, or device 116 (e.g., coming via a network)), while other
14 sources (e.g., storage device 112 or 114, or device 116) may simply be able to
15 transmit MIDI data that has been generated elsewhere.

16 In addition to being sources of MIDI data, devices 106 – 116 may also be
17 destinations for MIDI data. Some of the sources (e.g., keyboard 106, instruments
18 108, device 116, etc.) may be able to render (and possibly store) the audio data,
19 while other sources (e.g., storage devices 112 and 114) may only be able to store the
20 MIDI data.

21 The MIDI standard describes a technique for representing a musical piece
22 as a sequence of discrete notes and other events (e.g., such as might be performed
23 by an instrumentalist). These notes and events (the MIDI data) are communicated
24 in messages that are typically two or three bytes in length. These messages are
25 commonly classified as Channel Voice Messages, Channel Mode Messages, or

1 System Messages. Channel Voice Messages carry musical performance data
2 (corresponding to a specific channel), Channel Mode Messages affect the way a
3 receiving instrument will respond to the Channel Voice Messages, and System
4 Messages are control messages intended for all receivers in the system and are not
5 channel-specific. Examples of such messages include note on and note off
6 messages identifying particular notes to be turned on or off, aftertouch messages
7 (e.g., indicating how long a keyboard key has been held down after being pressed),
8 pitch wheel messages indicating how a pitch wheel has been adjusted, etc.
9 Additional information regarding the MIDI standard is available from the MIDI
10 Manufacturers Association of La Habra, California.

11 In the discussion herein, embodiments of the invention are described in the
12 general context of computer-executable instructions, such as program modules,
13 being executed by one or more conventional personal computers. Generally,
14 program modules include routines, programs, objects, components, data structures,
15 etc. that perform particular tasks or implement particular abstract data types.
16 Moreover, those skilled in the art will appreciate that various embodiments of the
17 invention may be practiced with other computer system configurations, including
18 hand-held devices, gaming consoles, Internet appliances, multiprocessor systems,
19 microprocessor-based or programmable consumer electronics, network PCs,
20 minicomputers, mainframe computers, and the like. In a distributed computer
21 environment, program modules may be located in both local and remote memory
22 storage devices.

23 Alternatively, embodiments of the invention can be implemented in
24 hardware or a combination of hardware, software, and/or firmware. For example,
25

1 at least part of the invention can be implemented in one or more application
2 specific integrated circuits (ASICs) or programmable logic devices (PLDs).

3 Fig. 2 shows a general example of a computer 142 that can be used in
4 accordance with certain embodiments of the invention. Computer 142 is shown as
5 an example of a computer that can perform the functions of computing device 102
6 of Fig. 1.

7 Computer 142 includes one or more processors or processing units 144, a
8 system memory 146, and a bus 148 that couples various system components
9 including the system memory 146 to processors 144. The bus 148 represents one
10 or more of any of several types of bus structures, including a memory bus or
11 memory controller, a peripheral bus, an accelerated graphics port, and a processor
12 or local bus using any of a variety of bus architectures. The system memory
13 includes read only memory (ROM) 150 and random access memory (RAM) 152.
14 A basic input/output system (BIOS) 154, containing the basic routines that help to
15 transfer information between elements within computer 142, such as during start-
16 up, is stored in ROM 150.

17 Computer 142 further includes a hard disk drive 156 for reading from and
18 writing to a hard disk, not shown, connected to bus 148 via a hard disk driver
19 interface 157 (e.g., a SCSI, ATA, or other type of interface); a magnetic disk drive
20 158 for reading from and writing to a removable magnetic disk 160, connected to
21 bus 148 via a magnetic disk drive interface 161; and an optical disk drive 162 for
22 reading from or writing to a removable optical disk 164 such as a CD ROM, DVD,
23 or other optical media, connected to bus 148 via an optical drive interface 165.
24 The drives and their associated computer-readable media provide nonvolatile
25 storage of computer readable instructions, data structures, program modules and

1 other data for computer 142. Although the exemplary environment described
2 herein employs a hard disk, a removable magnetic disk 160 and a removable
3 optical disk 164, it should be appreciated by those skilled in the art that other types
4 of computer readable media which can store data that is accessible by a computer,
5 such as magnetic cassettes, flash memory cards, digital video disks, random access
6 memories (RAMs) read only memories (ROM), and the like, may also be used in
7 the exemplary operating environment.

8 A number of program modules may be stored on the hard disk, magnetic
9 disk 160, optical disk 164, ROM 150, or RAM 152, including an operating system
10 170, one or more application programs 172, other program modules 174, and
11 program data 176. A user may enter commands and information into computer
12 142 through input devices such as keyboard 178 and pointing device 180. Other
13 input devices (not shown) may include a microphone, joystick, game pad, satellite
14 dish, scanner, or the like. These and other input devices are connected to the
15 processing unit 144 through an interface 168 that is coupled to the system bus. A
16 monitor 184 or other type of display device is also connected to the system bus
17 148 via an interface, such as a video adapter 186. In addition to the monitor,
18 personal computers typically include other peripheral output devices (not shown)
19 such as speakers and printers.

20 Computer 142 optionally operates in a networked environment using
21 logical connections to one or more remote computers, such as a remote computer
22 188. The remote computer 188 may be another personal computer, a server, a
23 router, a network PC, a peer device or other common network node, and typically
24 includes many or all of the elements described above relative to computer 142,
25 although only a memory storage device 190 has been illustrated in Fig. 2. The

logical connections depicted in Fig. 2 include a local area network (LAN) 192 and a wide area network (WAN) 194. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets, and the Internet. In the described embodiment of the invention, remote computer 188 executes an Internet Web browser program (which may optionally be integrated into the operating system 170) such as the "Internet Explorer" Web browser manufactured and distributed by Microsoft Corporation of Redmond, Washington.

When used in a LAN networking environment, computer 142 is connected to the local network 192 through a network interface or adapter 196. When used in a WAN networking environment, computer 142 typically includes a modem 198 or other component for establishing communications over the wide area network 194, such as the Internet. The modem 198, which may be internal or external, is connected to the system bus 148 via an interface (e.g., a serial port interface 168). In a networked environment, program modules depicted relative to the personal computer 142, or portions thereof, may be stored in the remote memory storage device. It is to be appreciated that the network connections shown are exemplary and other means of establishing a communications link between the computers may be used.

Computer 142 also optionally includes one or more broadcast tuners 200. Broadcast tuner 200 receives broadcast signals either directly (e.g., analog or digital cable transmissions fed directly into tuner 200) or via a reception device (e.g., via antenna 110 or satellite dish 114 of Fig. 1).

Generally, the data processors of computer 142 are programmed by means of instructions stored at different times in the various computer-readable storage media of the computer. Programs and operating systems are typically distributed,

for example, on floppy disks or CD-ROMs. From there, they are installed or loaded into the secondary memory of a computer. At execution, they are loaded at least partially into the computer's primary electronic memory. The invention described herein includes these and other various types of computer-readable storage media when such media contain instructions or programs for implementing the steps described below in conjunction with a microprocessor or other data processor. The invention also includes the computer itself when programmed according to the methods and techniques described below. Furthermore, certain sub-components of the computer may be programmed to perform the functions and steps described below. The invention includes such sub-components when they are programmed as described. In addition, the invention described herein includes data structures, described below, as embodied on various types of memory media.

For purposes of illustration, programs and other executable program components such as the operating system are illustrated herein as discrete blocks, although it is recognized that such programs and components reside at various times in different storage components of the computer, and are executed by the data processor(s) of the computer.

Kernel-Mode Processing

Fig. 3 is a block diagram illustrating an exemplary MIDI processing architecture in accordance with certain embodiments of the invention. The architecture 308 includes application(s) 310, graph builder 312, a MIDI transform module graph 314, and hardware devices 316 and 318. Hardware devices 316 and 318 are intended to represent any of a wide variety of MIDI data input and/or

1 output devices, such as any of devices 104 – 116 of Fig. 1. Hardware devices 316
2 and 318 are implemented in hardware level 320 of architecture 308.

3 Hardware devices 316 and 318 communicate with MIDI transform module
4 graph 314, passing input data to modules in graph 314 and receiving data from
5 modules in graph 314. Hardware devices 316 and 318 communicate with modules
6 in MIDI transform module graph 314 via hardware (HW) drivers 322 and 324,
7 respectively. A portion of each of hardware drivers 322 and 324 is implemented
8 as a module in graph 314 (these portions are often referred to as "miniport
9 streams"), and a portion is implemented in software external to graph 314 (often
10 referred to as "miniport drivers"). For input of MIDI data from a hardware device
11 316 (or 318), the hardware driver 322 (or 324) reads the data off of the hardware
12 device 316 (or 318) and puts the data in a form expected by the modules in graph
13 314. For output of MIDI data to a hardware device 316 (or 318), the hardware
14 driver receives the data and writes this data to the hardware directly.

15 An additional "feeder" module may also be included that is situated
16 between the miniport stream and the rest of the graph 314. Such feeder modules
17 are particularly useful in situations where the miniport driver is not aware of the
18 graph 314 or the data formats and protocols used within graph 314. In such
19 situations, the feeder module operates to convert formats between the hardware
20 (and hardware driver) specific format and the format supported by graph 314.
21 Essentially, for older miniport drivers whose miniport streams don't communicate
22 in the format supported by graph 314, the FeederIn and FeederOut modules
23 function as their liaison into that graph.

24 MIDI transform module graph 314 includes multiple (*n*) modules 326 (also
25 referred to as filters or MXFs (MIDI transform filters)) that can be coupled

1 together. Different source to destination paths (e.g., hardware device to hardware
2 device, hardware device to application, application to hardware device, etc.) can
3 exist within graph 314, using different modules 326 or sharing modules 326. Each
4 module 326 performs a particular function in processing MIDI data. Examples of
5 modules 326 include a sequencer to control the output of MIDI data to hardware
6 device 316 or 318 for playback, a packer module to package MIDI data for output
7 to application 310, etc. The operation of modules 326 is discussed in further detail
8 below.

9 Modern operating systems (e.g., those in the Microsoft Windows® family
10 of operating systems) typically include multiple privilege levels, often referred to
11 as user and kernel modes of operation (also called “ring 3” and “ring 0”). Kernel-
12 mode is usually associated with and reserved for portions of the operating system.
13 Kernel-mode (or “ring 0”) components run in a reserved address space, which is
14 protected from user-mode components. User-mode (or “ring 3”) components have
15 their own respective address spaces, and can make calls to kernel-mode
16 components using special procedures that require so-called “ring transitions” from
17 one privilege level to another. A ring transition involves a change in execution
18 context, which involves not only a change in address spaces, but also a transition
19 to a new processor state (including register values, stacks, privilege mode, etc).
20 As discussed above, such ring transitions can result in considerable latency and an
21 unpredictable amount of time.

22 MIDI transform module graph 314 is implemented in kernel-mode of
23 software level 328. Modules 326 are all implemented in kernel-mode, so no ring
24 transitions are required during the processing of MIDI data. Modules 326 are
25 implemented at a deferred procedure call (DPC) level, such as

1 DISPATCH_LEVEL. By implementing modules 326 at a higher priority level
2 than other user-mode software components, the modules 326 will have priority
3 over the user-mode components, thereby reducing delays in executing modules
4 326 and thus reducing latency and unpredictability in the transmitting and
5 processing of MIDI data.

6 In the illustrated example, modules 326 are implemented using Win32®
7 Driver Model (WDM) Kernel Streaming filters, thereby reducing the amount of
8 overhead necessary in communicating between modules 326. A low-overhead
9 interface is used by modules 326 to communicate with one another, rather than
10 higher-overhead I/O Request Packets (IRPs), and is described in more detail
11 below. Additional information regarding the WDM Kernel Streaming architecture
12 is available from Microsoft Corporation of Redmond, Washington.

13 Software level 328 also includes application(s) 310 implemented in user-
14 mode, and graph builder 312 implemented in kernel-mode. Any number of
15 applications 310 can interface with graph 314 (concurrently, in the event of a
16 multi-tasking operating system). Application 310 represents any of a wide variety
17 of applications that may use MIDI data. Examples of such applications include
18 games, reference materials (e.g., dictionaries or encyclopedias) and audio
19 programs (e.g., audio player, audio mixer, etc.).

20 In the illustrated example, graph builder 312 is responsible for generating a
21 particular graph 314. MIDI transform module graph 314 can vary depending on
22 what MIDI processing is desired. For example, a pitch modification module 326
23 would be included in graph 314 if pitch modification is desired, but otherwise
24 would not be included. MIDI transform module graph 314 has multiple different
25 modules available to it, although only selected modules may be incorporated into

graph 314 at any particular time. In the illustrated example, MIDI transform module graph 314 can include multiple modules 326 that do not have connections to other modules 326 – they simply do not operate on received MIDI data. Alternatively, only modules that operate on received MIDI data may be included in graph 314, with graph builder 312 accessing a module library 330 to copy modules into graph 314 when needed.

In one implementation, graph builder 312 accesses one or more locations to identify which modules are available to it. By way of example, a system registry may identify the modules or an index associated with module library 330 may identify the modules. Whenever a new module is added to the system, an identification of the module is added to these one or more locations. The identification may also include a descriptor, usable by graph builder 312 and/or an application 310, to identify the type of functionality provided by the module.

Graph builder 312 communicates with the individual modules 326 to configure graph 314 to carry out the desired MIDI processing functionality, as indicated to graph builder 312 by application 310. Although illustrated as a separate application that is accessed by other user-mode applications (e.g., application 310), graph builder 312 may alternatively be implemented as part of another application (e.g., part of application 310), or may be implemented as a separate application or system process in user-mode.

Application 310 can determine what functionality should be included in MIDI transform module graph 314 (and thus what modules graph builder 312 should include in graph 314) in any of a wide variety of manners. By way of example, application 310 may provide an interface to a user (e.g., a graphical user interface) that allows the user to identify various alterations he or she would like

1 made to a musical piece. By way of another example, application 310 may be pre-
2 programmed with particular functionality of what alterations should be made to a
3 musical piece, or may access another location (e.g., a remote server computer) to
4 obtain the information regarding what alterations should be made to the musical
5 piece. Additionally, graph builder 312 may automatically insert certain
6 functionality into the graph, as discussed in more detail below.

7 Graph builder 312 can change the connections in MIDI transform module
8 graph 314 during operation of the graph. In one implementation, graph builder
9 312 pauses or stops operation of graph 314 temporarily in order to make the
10 necessary changes, and then resumes operation of the graph. Alternatively, graph
11 builder 312 may change connections in the graph without stopping its operation.
12 Graph builder 312 and the manner in which it manages graph 314 are discussed in
13 further detail below.

14 MIDI transform module graphs are thus readily extensible. Graph builder
15 312 can re-arrange the graph in any of a wide variety of manners to accommodate
16 the desires of an application 310. New modules can be incorporated into a graph
17 to process MIDI data, modules can be removed from the graph so they no longer
18 process MIDI data, connections between modules can be modified so that modules
19 pass MIDI data to different modules, etc.

20 Communication between applications 310 and MIDI transform module
21 graph 314 transitions between different rings, so some latency and temporal
22 unpredictability may be experienced. In one implementation, communication
23 between applications 310 (or graph builder 312) and a module 326 is performed
24 using conventional IRPs. However, the processing of the MIDI data is being
25

1 carried out in kernel-mode, so such latency and/or temporal unpredictability does
2 not adversely affect the processing of the MIDI data.

3 Fig. 4 is a block diagram illustrating an exemplary module 326 in
4 accordance with certain embodiments of the invention. In the illustrated example,
5 each module 326 in graph 314 includes a processing portion 332 in which the
6 operation of the module 326 is carried out (and which varies by module). Each
7 module 326 also includes four interfaces: SetState 333, PutMessage 334,
8 ConnectOutput 335, and DisconnectOutput 336.

9 The SetState interface 333 allows the state of a module 326 to be set (e.g.,
10 by an application 310 or graph builder 312). In one implementation, valid states
11 include run, acquire, pause, and stop. The run state indicates that the module is to
12 run and perform its particular function. The acquire and pause states are
13 transitional states that can be used to assist in transitioning between the run and
14 stop states. The stop state indicates that the module is to stop running (it won't
15 accept any inputs or provide any outputs). When the SetState interface 333 is
16 called, one of the four valid states is included as a parameter by the calling
17 component.

18 The PutMessage interface 334 allows MIDI data to be input to a module
19 326. When the PutMessage interface 334 is called by another module, a pointer to
20 the MIDI data being passed (e.g., a data packet, as discussed in more detail below)
21 is included as a parameter, allowing the pointer to the MIDI data to be forwarded
22 to processing portion 332 for processing of the MIDI data. The PutMessage
23 interface 334 is called by another module 326, after it has finished processing the
24 MIDI data it received, and which passes the processed MIDI data to the next
25 module in the graph 314. After processing portion 332 finishes processing the

1 MIDI data, the PutMessage interface on the next module in the graph is called by
2 processing portion 332 to transfer the processed MIDI data to the connected
3 module 326 (the next module in the graph, as discussed below).

4 The ConnectOutput interface 335 allows a module 326 to be programmed
5 with the connected module (the next module in the graph). The ConnectOutput
6 interface is called by graph builder 312 to identify to the module where the output
7 of the module should be sent. When the ConnectOutput interface 335 is called, an
8 identifier (e.g., pointer to) the next module in the graph is included as a parameter
9 by the calling component. The default connected output is the allocator (discussed
10 in more detail below). In one implementation (called a "splitter" module), a
11 module 326 can be programmed with multiple connected modules (e.g., by
12 programming the module 326 with the PutMessage interfaces of each of the
13 multiple connected modules), allowing outputs to multiple "next" modules in the
14 graph. Conversely, multiple modules can point at a single "next" output module
15 (e.g., multiple modules may be programmed with the PutMessage interface of the
16 same "next" module).

17 The DisconnectOutput interface 336 allows a module 326 to be
18 disconnected from whatever module it was previously connected to (via the
19 ConnectOutput interface). The DisconnectOutput interface 336 is called by graph
20 builder 312 to have the module 326 reset to a default connected output (the
21 allocator). When the DisconnectOutput interface 336 is called, an identifier (e.g.,
22 pointer to) the module being disconnected from is included as a parameter by the
23 calling component. In one implementation, calling the ConnectOutput interface
24 335 or DisconnectOutput interface 336 with a parameter of NULL also
25 disconnects the "next" reference. Alternatively, the DisconnectOutput interface

1 336 may not be included (e.g., disconnecting the module can be accomplished by
2 calling ConnnectOutput 335 with a NULL parameter, or with an identification of
3 the allocator module as the next module).

4 Additional interfaces 337 may also be included on certain modules,
5 depending on the functions performed by the module. Two such additional
6 interfaces 337 are illustrated in Fig. 4: a SetParameters interface 338 and a
7 GetParameters interface 339. The SetParameters interface 338 allows a module
8 326 to receive various operational parameters set (e.g., from applications 310 or
9 graph builder 312), which are maintained as parameters 340. For example, a
10 module 326 that is to alter the pitch of a particular note(s) can be programmed, via
11 the SetParameters interface 338, with which note is to be altered and/or how much
12 the pitch is to be altered.

13 The GetParameters interface 339 allows coefficients (e.g., operational
14 parameters maintained as parameters 340) previously sent to the module, or any
15 other information the module may have been storing in a data section 341 (such as
16 MIDI jitter performance profiling data, number of events left in the allocator's free
17 memory pool, how much memory is currently allocated by the allocator, how
18 many messages have been enqueued by a sequencer module, a breakdown by
19 channel and/or channel group of what messages have been enqueued by the
20 sequencer module, etc), to be retrieved. The GetParameters interface 339 and
21 SetParameters interface 338 are typically called by graph builder 312, although
22 other applications 310 or modules in graph 314 could alternatively call them.

23 Returning to Fig. 3, one particular module that is included in MIDI
24 transform module graph 314 is referred to as the allocator. The allocator module
25 is responsible for obtaining memory from the memory manager (not shown) of the

computing device and making portions of the obtained memory available for MIDI data. The allocator module makes a pool of memory available for allocation to other modules in graph 314 as needed. The allocator module is called by another module 326 when MIDI data is received into the graph 314 (e.g., from hardware device 316 or 318, or application 310). The allocator module is also called when MIDI data is transferred out of the graph 314 (e.g., to hardware device 316 or 318, or application 310) so that memory that was being used by the MIDI data can be reclaimed and re-allocated for use by other MIDI data.

The allocator includes the interfaces discussed above, as well as additional interfaces that differ from the other modules 326. In the illustrated example, the allocator includes four additional interfaces: GetMessage, GetBufferSize, GetBuffer, and PutBuffer.

The GetMessage interface is called by another module 326 to obtain a data structure into which MIDI data can be input. The modules 326 communicate MIDI data to one another using a structure referred to as a data packet or event. Calling the GetMessage interface causes the allocator to return to the calling module a pointer to such a data packet in which the calling module can store MIDI data.

The PutMessage interface for the allocator takes a data structure and returns it to the free pool of packets that it maintains. This consists of its “processing.” The allocator is the original source and the ultimate destination of all event data structures of this type.

MIDI data is typically received in two or three byte messages. However, situations can arise where larger portions of MIDI data are received, referred to as System Exclusive, or SysEx messages. In such situations, the allocator allocates a

1 larger buffer for the MIDI data, such as 60 bytes or 4096 bytes. The
2 GetBufferSize interface is called by a module 326, and the allocator responds with
3 the size of the buffer that is (or will be) allocated for the portion of data. In one
4 implementation, the allocator always allocates buffers of the same size, so the
5 response by the allocator is always the same.

6 The GetBuffer interface is called by a module 326 and the allocator
7 responds by passing, to the module, a pointer to the buffer that can be used by the
8 module for the portion of MIDI data.

9 The PutBuffer interface is called by a module 326 to return the memory
10 space for the buffer to the allocator for re-allocation (the PutMessage interface
11 described above will call PutBuffer in turn, to return the memory space to the
12 allocator, if this hasn't been done already). When calling the PutBuffer interface,
13 the calling module includes, as a parameter, a pointer to the buffer being returned
14 to the allocator.

15 Situations can also arise where the amount of memory that is allocated by
16 the allocator for a buffer is smaller than the portion of MIDI data that is to be
17 received. In this situation, multiple buffers are requested from the allocator and
18 are "chained" together (e.g., a pointer in a data packet corresponding to each
19 identifies the starting point of the next buffer). An indication may also be made in
20 the corresponding data packet that identifies whether a particular buffer stores the
21 entire portion of MIDI data or only a sub-portion of the MIDI data.

22 Many modern processors and operating systems support virtual memory.
23 Virtual memory allows the operating system to allocate more memory to
24 application processes than is physically available in the computing device. Data
25 can then be swapped between physical memory (e.g., RAM) and another storage

device (e.g., a hard disk drive), a process referred to as paging. The use of virtual memory gives the appearance of more physical memory being available in the computing device than is actually available. The tradeoff, however, is that swapping data from a disk drive to memory typically takes significantly longer than simply retrieving the data directly from memory.

In one implementation, the allocator obtains non-pageable portions of memory from the memory manager. That is, the memory that is obtained by the allocator refers to a portion of physical memory that will not be swapped to disk. Thus, processing of MIDI data will not be adversely affected by delays in swapping data between memory and a disk.

In one implementation, each module 326, when added to graph 314, is passed an identifier (e.g., pointer to) the allocator module as well as a clock. The allocator module is used, as described above, to allow memory for MIDI data to be obtained and released. The clock is a common reference clock that is used by all of the modules 326 to maintain synchronization with one another. The manner in which the clock is used can vary, depending on the function performed by the modules. For example, a module may generate a time stamp, based on the clock, indicating when the MIDI data was received by the module, or may access a presentation time for the data indicating when it is to be played back.

Alternatively, some modules may not need, and thus need not include, pointers to the reference clock and/or the allocator module (however, in implementations where the default output destination for each module is an allocator module, then each module needs a pointer to the allocator in order to properly initialize). For example, if a module will carry out its functionality

1 without regard for what the current reference time is, then a pointer to the
2 reference clock is not necessary.

3 Fig. 5 is a block diagram illustrating an exemplary MIDI message 345.
4 MIDI message 345 includes a status portion 346 and a data portion 347. Status
5 portion 346 is one byte, while data portion 347 is either one or two bytes. The size
6 of data portion 347 is encoded in the status portion 346 (either directly, or
7 inherently based on some other value (such as the type of command)). The MIDI
8 data is received from and passed to hardware devices 316 and 318 of Fig. 3, and
9 possibly application 310, as messages 345. Typically each message 345 identifies
10 a single command (e.g., note on, note off, change volume, pitch bend, etc.). The
11 audio data included in data portion 347 will vary depending on the message type.

12 Fig. 6 is a block diagram illustrating an exemplary MIDI data packet 350 in
13 accordance with certain embodiments of the invention. MIDI data (or references,
14 such as pointers, thereto) is communicated among modules 326 in MIDI transform
15 module graph 314 of Fig. 3 as data packets 350, also referred to as events. When a
16 MIDI message 345 of Fig. 5 is received into graph 314, the receiving module 326
17 generates a data packet 350 that incorporates the message.

18 Data packet 350 includes a reserved portion 352 (e.g., one byte), a structure
19 byte count portion 354 (e.g., one byte), an event byte count portion 356 (e.g. two
20 bytes), a channel group portion 358 (e.g., two bytes), a flags portion 360 (e.g. two
21 bytes), a presentation time portion 362 (e.g., eight bytes), a byte position 364 (e.g.,
22 eight bytes), a next event portion 366 (e.g. four bytes), and a data portion 368
23 (e.g., four bytes). Reserved portion 352 is reserved for future use. Structure byte
24 count portion 354 identifies the size of the message 350.

1 Event byte count portion 356 identifies the number of data bytes that are
2 referred to in data portion 368. The number of data bytes could be the number
3 actually stored in data portion 368 (e.g., two or three, depending on the type of
4 MIDI data), or alternatively the number of bytes pointed to by a pointer in data
5 portion 368, (e.g., if the number of data bytes is greater than the size of a pointer).
6 If the event is a package event (pointing to a chain of events, as discussed in more
7 detail below), then the portion 356 has no value. Alternatively, portion 356 could
8 be set to the value of event byte count portion 356 of the first regular event in its
9 chain, or to the byte count of the entire long message. If event portion 356 is not
10 set to the byte count of the entire long message, then data could still be flowing
11 into the last message structure of the package event while the initial data is already
12 being processed elsewhere.

13 Channel group portion 358 identifies which of multiple channel groups the
14 data identified in data portion 368 corresponds to. The MIDI standard supports
15 sixteen different channels, allowing essentially sixteen different instruments or
16 "voices" to be processed and/or played concurrently for a musical piece. Use of
17 channel groups allows the number of channels to be expanded beyond sixteen.
18 Each channel group can refer to any one of sixteen channels (as encoded in status
19 byte 346 of message 345 of Fig. 5). In one implementation, channel group portion
20 358 is a 2-byte value, allowing up to 65,536 (64k) different channel groups to be
21 identified (as each channel group can have up to sixteen channels, this allows a
22 total of 1,048,576 (1Meg) different channels).

23 Flags portion 360 identifies various flags that can be set regarding the MIDI
24 data corresponding to data packet 350. In one implementation, zero or more of
25 multiple different flags can be set: an Event In Use (EIU) flag, an Event

1 Incomplete (EI) flag, one or more MIDI Parse State flags (MPS), or a Package
2 Event (PE) flag. The Event In Use flag should always be on (set) when an event is
3 traveling through the system; when it is in the free pool this bit should be cleared.
4 This is used to prevent memory corruption. The Event Incomplete flag is set if the
5 event continues beyond the buffer pointed to by data portion 368, or if the
6 message is a System Exclusive (SysEx) message. The MIDI Parse State flags are
7 used by a capture sink module (or other module parsing an unparsed stream of
8 MIDI data) in order to keep track of the state of the unparsed stream of MIDI data.
9 As the capture sink module successfully parses the MIDI data into a complete
10 message, these two bits should be cleared. In one implementation these flags have
11 been removed from the public flags field.

12 The Package Event flag is set if data packet 350 points to a chain of other
13 packets 350 that should be dealt with atomically. By way of example, if a portion
14 of MIDI data is being processed that is large enough to require a chain of data
15 packets 350, then this packet chain should be passed around atomically (e.g., not
16 separated so that a module receives only a portion of the chain). Setting the
17 Package Event flag identifies data field 374 as pointing to a chain of multiple
18 additional packets 350.

19 Presentation time portion 362 specifies the presentation time for the data
20 corresponding to data packet 350 (i.e., for an event). The presentation of an event
21 depends on the type of event: note on events are presented by rendering the
22 identified note, note off events are presented by ceasing rendering of the identified
23 note, pitch bend events are presented by altering the pitch of the identified note in
24 the identified manner, etc. A module 326 of Fig. 3, by comparing the current
25 reference clock time to the presentation time identified in portion 362, can

1 determine when, relative to the current time, the event should be presented to a
2 hardware device 316 or 318. In one implementation, portion 362 identifies
3 presentation times in 100 nanosecond (ns) units.

4 Byte position portion 364 identifies where this message (included in data
5 portion 368) is situated in the overall stream of bytes from the application (e.g.,
6 application 310 of Fig. 3). Because certain applications use the release of their
7 submitted buffers as a timing mechanism, it is important to keep track of how far
8 processing has gone in the byte order, and release buffers only up to that point
9 (and only release those buffers back to the application after the corresponding
10 bytes have actually been played). In this case the allocator module looks at the
11 byte offset when a message is destroyed (returned for re-allocation), and alerts a
12 stream object (e.g., the IRP stream object used to pass the buffer to graph 314) that
13 a certain amount of memory can be released up to the client application.

14 Next event portion 366 identifies the next packet 350 in a chain of packets,
15 if any. If there is no next packet, then next event portion 366 is NULL.

16 Data portion 368 can include one of three things: packet data 370 (a
17 message 345 of Fig. 5), a pointer 372 to a chain of packets 350, or a pointer 374 to
18 a data buffer. Which of these three things is included in data portion 368 can be
19 determined based on the value in event byte count field 356 and/or flags portion
20 360. In the illustrated example, the size of a pointer is greater than three bytes
21 (e.g., is 4 bytes). If the event byte count field 356 is less than or equal to the size
22 of a pointer, then data portion 368 includes packet data 370; otherwise data portion
23 368 includes a pointer 374 to a data buffer. However, this determination is
24 overridden if the Package Event flag of flags portion 360 is set, which indicates
25

1 that data portion 368 includes a pointer 372 to a chain of packets (regardless of the
2 value of event byte count field 356).

3 Returning to Fig. 3, certain modules 326 may receive MIDI data from
4 application 310 and/or send MIDI data to application 310. In the illustrated
5 example, MIDI data can be received from and/or sent to an application 310 in
6 different formats, depending at least in part on whether application 310 is aware of
7 the MIDI transform module graph 314 and the format of data packets 350 (of Fig.
8 5) used in graph 314. If application 310 is not aware of the format of data packets
9 350 then application 310 is referred to as a "legacy" application and the MIDI data
10 received from application 310 is converted into the format of data packets 350.
11 Application 310, whether a legacy application or not, communicates MIDI data to
12 (or receives MIDI data from) a module 326 in a buffer including one or more
13 MIDI messages (or data packets 350).

14 Fig. 7 is a block diagram illustrating an exemplary buffer for
15 communicating MIDI data between a non-legacy application and a MIDI
16 transform module graph module in accordance with certain embodiments of the
17 invention. A buffer 380, which can be used to store one or more packaged data
18 packets, is illustrated including multiple packaged data packets 382 and 384. Each
19 packaged data packet 382 and 384 includes a data packet 350 of Fig. 6 as well as
20 additional header information. This combination of data packet 350 and header
21 information is referred to as a packaged data packet. In one implementation,
22 packaged data packets are quadword (8-byte) aligned for alignment and speed
23 reasons (e.g., by adding padding 394 as needed).

24 The header information for each packaged data packet includes an event
25 byte count portion 386, a channel group portion 388, a reference time delta portion

1 390, and a flags portion 392. The event byte count portion 386 identifies the
2 number of bytes in the event(s) corresponding to data packet 350 (which is the
3 same value as maintained in event portion 356 of data packet 350 of Fig. 6, unless
4 the packet is broken up into multiple events structures.). The channel group
5 portion 388 identifies which of multiple channel groups the event(s) corresponding
6 to data packet 350 correspond to (which is the same value as maintained in
7 channel group portion 358 of data packet 350).

8 The reference time delta portion 390 identifies the difference in
9 presentation time between packaged data packet 382 (stored in presentation time
10 portion 362 of data packet 350 of Fig. 6) and the beginning of buffer 380. The
11 beginning time of buffer 380 can be identified as the presentation time of the first
12 packaged data packet 382 in buffer 380, or alternatively buffer 380 may have a
13 corresponding start time (based on the same reference clock as the presentation
14 time of data packets 350 are based on).

15 Flags portion 392 identifies one or more flags that can be set regarding the
16 corresponding data packet 350. In one implementation, only one flag is
17 implemented - an Event Structured flag that is set to indicate that structured data is
18 included in data packet 350. Structured data is expected to parse correctly from a
19 raw MIDI data stream into complete message packets. An unstructured data
20 stream is perhaps not MIDI compliant, so it isn't grouped into MIDI messages like
21 a structured stream is – the original groupings of bytes of unstructured data are
22 unmodified. Whether the data is compliant (structured) or non-compliant
23 (unstructured) is indicated by the Event Structured flag.

24 Fig. 8 is a block diagram illustrating an exemplary buffer for
25 communicating MIDI data between a legacy application and a MIDI transform

1 module graph module in accordance with certain embodiments of the invention.
2 A buffer 410, which can be used to store one or more packaged events, is
3 illustrated including multiple packaged events 412 and 414. Each packaged event
4 412 and 414 includes a message 345 of Fig. 5 as well as additional header
5 information. This combination of message 345 and header information is referred
6 to as a packaged event (or packaged message). In one implementation, packaged
7 events are quadword (8-byte) aligned for speed and alignment reasons (e.g., by
8 adding padding 420 as needed).

9 The additional header information in each packaged event includes a time
10 delta portion 416 and a byte count portion 418. Time delta portion 416 identifies
11 the difference between the presentation time of the packaged event and the
12 presentation time of the immediately preceding packaged event. These
13 presentation times are established by the legacy application passing the MIDI data
14 to the graph. For the first packaged event in buffer 410, time delta portion 416
15 identifies the difference between the presentation time of the packed event and the
16 beginning time corresponding to buffer 410. The beginning time corresponding to
17 buffer 410 is the presentation time for the entire buffer (the first message in the
18 buffer can have some positive offset in time and does not have to start right at the
19 head of the buffer).

20 Byte count portion 416 identifies the number of bytes in message 345.

21 Fig. 9 is a block diagram illustrating an exemplary MIDI transform module
22 graph 430 such as may be used in accordance with certain embodiments of the
23 invention. In the illustrated example, keys on a keyboard can be activated and the
24 resultant MIDI data forwarded to an application executing in user-mode as well as
25

1 being immediately played back. Additionally, MIDI data can be input to graph
2 430 from a user-mode application for playback.

3 One source of MIDI data in Fig. 9 is keyboard 432, which provides the
4 MIDI data as a raw stream of MIDI bytes via a hardware driver including a
5 miniport stream (in) module 434. Module 434 calls the GetMessage interface of
6 allocator 436 for memory space (a data packet 350) into which a structured packet
7 can be placed, and module 434 adds a timestamp to the data packet 350.
8 Alternatively, module 434 may rely on capture sink module 438, discussed below,
9 to generate the packets 350, in which case module 434 adds a timestamp to each
10 byte of the raw data it receives prior to forwarding the data to capture sink module
11 438. In the illustrated example, notes are to be played immediately upon
12 activation of the corresponding key on keyboard 432, so the timestamp stored by
13 module 434 as the presentation time of the data packets 350 is the current reading
14 of the master (reference) clock.

15 Module 434 is connected to capture sink module 438, splitter module 430
16 or packer 442 (the splitter module is optional – only inserted if, for example, the
17 graph builder has been told to connect "kernel THRU"). Capture sink module 438
18 is optional, and operates to generate packets 350 from a received MIDI data byte
19 stream. If module 434 generates packets 350, then capture sink 438 is not
20 necessary and module 434 is connected to optional splitter module 440 or packer
21 442. However, if module 434 does not generate packets 350, then module 434 is
22 connected to capture sink module 438. After adding the timestamp, module 434
23 calls the PutMessage interface of the module it is connected to (either capture sink
24 module 438, splitter module 440 or packer module 442), which passes the newly
25 created message to that module.

1 The manner in which packets 350 are generated from the received raw
2 MIDI data byte stream (regardless of whether it is performed by module 434 or
3 capture sink module 438) is dependent on the particular type of data (e.g., the data
4 may be included in data portion 368 (Fig. 6), a pointer may be included in data
5 portion 368, etc.). In situations where multiple bytes of raw MIDI data are being
6 stored in data portion 368, the timestamp of the first of the multiple bytes is used
7 as the timestamp for the packet 350. Additionally, situations can arise where
8 additional event structures have been obtained from allocator 436 than are actually
9 needed (e.g., multiple bytes were not received together and multiple event
10 structures were received for each, but they are to be grouped together in the same
11 event structure). In such situations the additional event structures can be kept for
12 future MIDI data, or alternatively returned to allocator 436 for re-allocation.

13 Splitter module 440 operates to duplicate received data packets 350 and
14 forward each to a different module. In the illustrated example, splitter module 440
15 is connected to both packer module 442 and sequencer module 444. Upon receipt
16 of a data packet 350, splitter module 440 obtains additional memory space from
17 allocator 436, copies the contents of the received packet into the new packet
18 memory space, and calls the PutMessage interfaces of the modules it is connected
19 to, which passes one data packet 350 to each of the connected modules (i.e., one
20 data packet to packer module 442 and one data packet to sequencer module 444).
21 Splitter module 440 may optionally operate to duplicate a received data packet 350
22 only if the received data packet corresponds to audio data matching a particular
23 type, such as certain note(s), channel(s), and/or channel group(s).

24 Packer module 442 operates to combine one or more received packets into
25 a buffer (such as buffer 380 of Fig. 7 or buffer 410 of Fig. 8) and forward the

1 buffer to a user-mode application (e.g., using IRPs with a message format desired
2 by the application). Two different packer modules can be used as packer module
3 442, one being dedicated to legacy applications and the other being dedicated to
4 non-legacy applications. Alternatively, a single packer module may be used and
5 the type of buffer (e.g., buffer 380 or 410) used by packer module 442 being
6 dependent on whether the application to receive the buffer is a legacy application.

7 Once a data packet is forwarded to the user-mode application, packer 442
8 calls its programmed PutMessage interface (the PutMessage interface that the
9 module packer 442 is connected to) for that packet. Packer module 442 is
10 connected to allocator module 436, so calling its programmed PutMessage
11 interface for a data packet returns the memory space used by the data packet to
12 allocator 436 for re-allocation. Alternatively, packer 442 may wait to call
13 allocator 436 for each packet in the buffer after the entire buffer is forwarded to
14 the user-mode application.

15 Sequencer module 444 operates to control the delivery of data packets 350
16 received from splitter module 440 to miniport stream (out) module 446 for playing
17 on speakers 450. Sequencer module 444 does not change the data itself, but
18 module 444 does reorder the data packets by timestamp and delay the calling of
19 PutMessage (to forward the message on) until the appropriate time. Sequencer
20 module 444 is connected to module 446, so calling PutMessage causes sequencer
21 module 444 to forward a data packet to module 446. Sequencer module 444
22 compares the presentation times of received data packets 350 to the current
23 reference time. If the presentation time is equal to or earlier than the current time
24 then the data packet 350 is to be played back immediately and the PutMessage
25 interface is called for the packet. However, if the presentation time is later than

1 the current time, then the data packet 350 is queued until the presentation time is
2 equal to the current time, at which point sequencer module 444 calls its
3 programmed PutMessage interface for the packet. In one implementation,
4 sequencer 444 is a high-resolution sequencer, measuring time in 100 ns units.

5 Alternatively, sequencer module 444 may attempt to forward packets to
6 module 446 slightly in advance of their presentation time (that is, when the
7 presentation time of the packet is within a threshold amount of time later than the
8 current time). The amount of this threshold time would be, for example, an
9 anticipated amount of time that is necessary for the data packet to pass through
10 module 446 and to speakers 450 for playing, resulting in playback of the data
11 packets at their presentation times rather than submission of the packets to module
12 446 at their presentation times. An additional "buffer" amount of time may also be
13 added to the anticipated amount of time to allow output module 448 (or speakers
14 450) to have the audio messages delivered at a particular time (e.g., five seconds
15 before the data needs to be rendered by speakers 450).

16 A module 446 could furthermore specify that it did not want the sequencer
17 to hold back the data at all, even if data were extremely early. In this case, the
18 HW driver "wants to do its own sequencing," so the sequencer uses a very high
19 threshold (or alternatively a sequencer need not be inserted above this particular
20 module 446). The module 446 is receiving events with presentation timestamps in
21 them, and it also has access to the clock (e.g., being handed a pointer to it when it
22 was initialized), so if the module 446 wanted to synchronize that clock to its own
23 very-high performance clock (such as an audio sample clock), it could potentially
24 achieve even higher resolution and lower jitter than the built-in clock/sequencer.
25

Module 446 operates as a hardware driver customized to the MIDI output device 450. Module 446 converts the information in the received data packets 350 to a form specific to the output device 450. Different manufacturers can use different signaling techniques, so the exact manner in which module 446 operates will vary based on speakers 450 (and/or output module 448). Module 446 is coupled to an output module 448 which synthesizes the MIDI data into sound that can be played by speakers 450. Although illustrated in the software level, output module 448 may alternatively be implemented in the hardware level. By way of example, module 446 may be a MIDI output module which synthesizes MIDI messages into sound, a MIDI-to-waveform converter (often referred to as a software synthesizer), etc. In one implementation, output module 448 is included as part of a hardware driver corresponding to output device 450.

Module 446 is connected to allocator module 436. After the data for a data packet has been communicated to the output device 450, module 446 calls the PutMessage interface of the module it is connected to (allocator 436) to return the memory space used by the data packet to allocator 436 for re-allocation.

Another source of MIDI data illustrated in Fig. 9 is a user-mode application(s). A user-mode application can transmit MIDI data to unpacker module 452 in a buffer (such as buffer 380 of Fig. 7 or buffer 410 of Fig. 8). Analogous to packer module 442 discussed above, different unpacker modules can be used as unpacker module 452, (one being dedicated to legacy applications and the other being dedicated to non-legacy applications), or alternatively a single dual-mode unpacker module may be used. Unpacker module 452 operates to convert the MIDI data in the received buffer into data packets 350, obtaining memory space from allocator module 436 for generation of the data packets 350.

1 Unpacker module 452 is connected to sequencer module 444. Once a data packet
2 350 is created, unpacker module 452 calls its programmed PutMessage interface to
3 transmit the data packet 350 to sequencer module 444. Sequencer module 444,
4 upon receipt of the data packet 350, operates as discussed above to either queue
5 the data packet 350 or immediately transfer the data packet 350 to module 446.
6 Because the unpacker 450 has done its job of converting the data stream from a
7 large buffer into smaller individual data packets, these data packets can be easily
8 sorted and interleaved with a data stream also entering the sequencer 444 – from
9 the splitter 440 for example.

10 Fig. 10 is a block diagram illustrating another exemplary MIDI transform
11 module graph 454 such as may be used in accordance with certain embodiments of
12 the invention. Graph 454 of Fig. 10 is similar to graph 430 of Fig. 9, except that
13 one or more additional modules 456 that perform various operations are added to
14 graph 454 by graph builder 312 of Fig. 3. As illustrated, one or more of these
15 additional modules 456 can be added in graph 454 in a variety of different
16 locations, such as between modules 438 and 440, between modules 440 and 442,
17 between modules 440 and 444, between modules 452 and 444, and/or between
18 modules 444 and 446.

19 Fig. 11 is a flowchart illustrating an exemplary process for the operation of
20 a module in a MIDI transform module graph in accordance with certain
21 embodiments of the invention. In the illustrated example, the process of Fig. 11 is
22 implemented by a software module (e.g., module 326 of Fig. 3) executing on a
23 computing device.

24 Initially, a data packet including MIDI data (e.g., a data packet 350 of Fig.
25 5) is received by the module (act 462) (when its own PutMessage interface is

1 called). Upon receipt of the MIDI data, the module processes the MIDI data (act
2 464). The exact manner in which the data is processed is dependent on the
3 particular module, as discussed above. Once processing is complete, the
4 programmed PutMessage interface (which is on a different module) is called (act
5 468), forwarding the data packet to the next module in the graph.

6 Fig. 12 is a flowchart illustrating an exemplary process for the operation of
7 a graph builder in accordance with certain embodiments of the invention. In the
8 illustrated example, the process of Fig. 12 is carried out by a graph builder 312 of
9 Fig. 3 implemented in software. Fig. 12 is discussed with additional reference to
10 Fig. 3. Although a specific ordering of acts is illustrated in Fig. 12, the ordering of
11 the acts can alternatively be re-arranged.

12 Initially, graph builder 312 receives a request to build a graph (act 472).
13 This request may be for a new graph or alternatively to modify a currently existing
14 graph. The user-mode application 310 that submits the request to build the graph
15 includes an identification of the functionality that the graph should include. This
16 functionality can include any of a wide variety operations, including pitch bends,
17 volume changes, aftertouch alterations, etc. The user-mode application also
18 submits, if relevant, an ordering to the changes. By way of example, the
19 application may indicate that the pitch bend should occur prior to or subsequent to
20 some other alteration.

21 In response to the received request, graph builder 312 determines which
22 graph modules are to be included based at least in part on the desired functionality
23 identified in the request (act 474). Graph builder 312 is programmed with, or
24 otherwise has access to, information identifying which modules correspond to
25 which functionality. By way of example, a lookup table may be used that maps

1 functionality to module identifiers. Graph builder 312 also automatically adds
2 certain modules into the graph (if not already present). In one implementation, an
3 allocator module is automatically inserted, an unpacker module is automatically
4 inserted for each output path, and packer and capture sink modules are
5 automatically inserted for each input path.

6 Graph builder 312 also determines the connections among the graph
7 modules based at least in part on the desired functionality (and ordering, if any)
8 included in the request (act 476). In one implementation, graph builder 312 is
9 programmed with a set of rules regarding the building of graphs (e.g., which
10 modules must or should, if possible, be prior to which other modules in the graph).
11 Based on such a set of rules, the MIDI transform module graph can be constructed.

12 Graph builder 312 then initializes any needed graph modules (act 478).
13 The manner in which graph modules are initialized can vary depending on the type
14 of module. For example, pointers to the allocator module and reference clock may
15 be passed to the module, other operating parameters may be passed to the module,
16 etc.

17 Graph builder then adds any needed graph modules (as determined in act
18 474) to the graph (act 480), and connects the graph modules using the connections
19 determined in act 476 (act 482). If any modules need to be temporarily paused to
20 perform the connections, graph builder 312 changes the state of such graph
21 modules to a stop state (act 484). The outputs for the added modules are
22 connected first, and then the other modules are redirected to feed them, working in
23 a direction "up" the graph from destination to source (act 486). This reduces the
24 chances that the graph would need to be stopped to insert modules. Once
25 connected, any modules in the graph that are not already in a run state are started

(e.g., set to a run state) (act 488). Alternatively, another component may set the modules in the graph to the run state, such as application 310. In one implementation, the component (e.g., graph builder 312) setting the nodes in the graph to the run state follows a particular ordering. By way of example, the component may begin setting modules to run state at a MIDI data source and follow that through to a destination, then repeat for additional paths in the graph (e.g., in graph 430 of Fig. 8, the starting of modules may be in the following order: modules 436, 434, 438, 440, 442, 444, 446, 452). Alternatively, certain modules may be in a "start first" category (e.g., allocator 436 and sequencer 444 of Fig. 8).

In one implementation, graph builder 312 follows certain rules when adding or deleting items from the graph as well as when starting or stopping the graph. Reference is made herein to "merger" modules, branching modules, and branches within a graph. Merging is built-in to the interface described above, and a merger module refers to any module that has two or more other modules outputting to it (that is, two or more other modules calling its PutMessage interface). Graph builder 312 knows this information (who the mergers are), however the mergers themselves do not. A branching module refers to any module from which two or more branches extend (that is, any module that duplicates (at least in part) data and forwards copies of the data to multiple modules). An example of a branching module is a splitter module. A branch refers to a string of modules leading to or from (but not including) a branching module or merger module, as well as a string of modules between (but not including) merger and branching modules.

When moving the graph from a lower state (e.g., stop) to a higher state (e.g., run), graph builder 312 first changes the state of the destination modules,

1 then works its way toward the source modules. At places where the graph
2 branches (e.g., splitter modules), all destination branches are changed before the
3 branching module (e.g., splitter module) is changed. In this way, by the time the
4 "spigot is turned on" at the source, the rest of the graph is in run state and ready to
5 go.

6 When moving the graph from a higher state (e.g., run) to a lower state (e.g.,
7 stop), the opposite tack is taken. First graph builder 312 stops the source(s), then
8 continues stopping the modules as it progresses toward the destination module(s).
9 In this way the "spigot is turned off" at the source(s) first, and the rest of the graph
10 is given time for data to empty out and for the modules to "quiet" themselves. A
11 module quieting itself refers to any residual data in the module being emptied out
12 (e.g., an echo is passively allowed to die off, etc.). Quieting a module can also be
13 actively accomplished by putting the running module into a lower state (e.g., the
14 pause state) until it is no longer processing any residual data (which graph builder
15 312 can determine, for example, by calling its GetParameters interface).

16 When a module is in stop state, the module fails any calls to the module's
17 PutMessage interface. When the module is in the acquire state, the module
18 accepts PutMessage calls without failing them, but it does not forward messages
19 onward. When the module is in the pause state, it accepts PutMessage calls and
20 can work normally as long as it does not require the clock (if it needs a clock, then
21 the pause state is treated the same as the acquire state). Clockless modules are
22 considered "passive" modules that can operate fully during the "priming" sequence
23 when the graph is in the pause state. Active modules only operate when in the run
24 state. By way of example, splitter modules are passive, while sequencer modules,
25 miniport streams, packer modules, and unpacker modules are active.

Different portions of a graph can be in different states. When a source is inactive, all modules on that same branch can be inactive as well. Generally, all the modules in a particular branch should be in the same state, including source and destination modules if they are on that branch. Typically, the splitter module is put in the same state as its input module. A merger module is put in the highest state (e.g., in the order stop, pause, acquire, run) of any of its input modules.

Graph builder 312 can insert modules to or delete modules from a graph "live" (while the graph is running). In one implementation, any module except miniport streams, packers, unpackers, capture sinks, and sequencers can be inserted to or deleted from the graph while the graph is running. If a module is to be added or deleted while the graph is running, care should be taken to ensure that no data is lost when making changes, and when deleting a module that the module is allowed to completely quiet itself before it is disconnected.

By way of example, when adding a module B between modules A and C, first the output of module B is connected to the input of module C (module C is still being fed by module A). Then, graph builder 312 switches the output of module A from module C to module B with a single ConnectOutput call. The module synchronizes ConnectOutput calls with PutMessage calls, so accomplishing the graph change with a single ConnectOutput call ensures that no data packets are lost during the switchover. In the case of a branching module, all of its outputs are connected first, then its source is connected. When adding a module immediately previous to a merger module (where the additional module is intended to be common to both data paths), the additional module becomes the new merger module, and the item that was previously considered a merger module is no longer regarded as a merger module. In that case, the new merger module's

1 output and the old merger module's input are connected first, then the old merger
2 module's inputs are switched to the new merger module's inputs. If it is absolutely
3 necessary that all of the merger module's inputs switch to the new merger at the
4 same instant, then a special SetParams call should be made to each of the
5 "upstream" input modules to set a timestamp for when the ConnectOutput should
6 take place.

7 When deleting a module B from between modules A and C, first the output
8 of module A is connected to the input of module C (module B is effectively
9 bypassed at this time). Then, after module B empties and quiets itself (e.g., it
10 might be an echo or other time-based effect), its output is reset to the allocator.
11 Then module B can be safely destroyed (e.g., removed from the graph). When
12 deleting a merger module, first its inputs are switched to the subsequent module
13 (which becomes a merger module now), then after the old merger module quiets,
14 its output is disconnected. When deleting a branching module, this is because an
15 entire branch is no longer needed. In that case, the branching module output going
16 to that branch is disconnected. If the branching module had more than two
17 outputs, then the graph builder calls DisconnectOutput to disconnect that output
18 from the branching module's output list. At that point the subsequent modules in
19 that branch can be safely destroyed. However, if the branching module had only
20 two connected outputs, then the splitter module is no longer necessary. In that
21 case, the splitter module is bypassed (the previous module's output is connected to
22 the subsequent module's input), then after the splitter module quiets it is
23 disconnected and destroyed.

Additional Transform Modules

Specific examples of modules that can be included in a MIDI transform module graph (such as graph 430 of Fig. 9, graph 454 of Fig. 10, or graph 314 of Fig. 3) are described above. Various additional modules can also be included in a MIDI transform module graph, allowing user-mode applications to generate any of a wide variety of audio effects. Furthermore, as graph builder 312 of Fig. 3 allows the MIDI transform module graph to be readily changed, the functionality of the MIDI transform module graph can be changed to include new modules as they are developed. Examples of additional modules that can be included in a MIDI transform module graph are described below.

Unpacker Modules. Unpacker modules, in addition to those discussed above, can also be included in a MIDI transform module graph. Unpacker modules operate to receive data into the graph from a user-mode application, converting the MIDI data received in the user-mode application format into data packets 350 (Fig. 6) for communicating to other modules in the graph. Additional unpacker modules, supporting any of a wide variety of user-mode application specific formats, can be included in the graph.

Packer Modules. Packer modules, in addition to those discussed above, can also be included in a MIDI transform module graph. Packer modules operate to output MIDI data from the graph to a user-mode application, converting the MIDI data from the data packets 350 into a user-mode application specific format. Additional packer modules, supporting any of a wide variety of user-mode application specific formats, can be included in the graph.

Feeder In Modules. A Feeder In module operates to convert MIDI data received in from a software component that is not aware of the data formats and

1 protocols used in a module graph (e.g., graph 314 of Fig. 3) into data packets 350.
2 Such components are typically referred to as "legacy" components, and include,
3 for example, older hardware miniport drivers. Different Feeder In modules can be
4 used that are specific to the particular hardware drivers they are receiving the
5 MIDI data from. The exact manner in which the Feeder In modules operate will
6 vary, depending on what actions are necessary to convert the received MIDI data
7 to the data packets 350.

8 Feeder Out Modules. A Feeder Out module operates to convert MIDI data
9 in data packets 350 into the format expected by a particular legacy component
10 (e.g., older hardware miniport driver) that is not aware of the data formats and
11 protocols used in a module graph (e.g., graph 314 of Fig. 3). Different Feeder Out
12 modules can be used that are specific to the particular hardware drivers they are
13 sending the MIDI data to. The exact manner in which the Feeder Out modules
14 operate will vary, depending on what actions are necessary to convert the MIDI
15 data in the data packets 350 into the format expected by the corresponding
16 hardware driver.

17

18 **Conclusion**

19 Although the description above uses language that is specific to structural
20 features and/or methodological acts, it is to be understood that the invention
21 defined in the appended claims is not limited to the specific features or acts
22 described. Rather, the specific features and acts are disclosed as exemplary forms
23 of implementing the invention.

1 **CLAIMS**

2

3 1. An extensible architecture for kernel-mode processing of audio
4 messages, the architecture comprising:

5 a plurality of modules coupled together, in a module graph implemented in
6 kernel-mode, each of the modules performing an operation in the processing of the
7 audio messages; and

8 a plurality of additional modules that can be added to the module graph,
9 wherein each of the plurality of additional modules is coupled to one of the
10 plurality of modules.

11

12 2. The extensible architecture of claim 1, wherein the audio messages
13 include audio data.

14

15 3. The extensible architecture of claim 1, wherein the audio messages
16 include control information.

17

18 4. The extensible architecture of claim 1, further comprising a graph
19 builder component, communicatively coupled to the plurality of modules, to add
20 selected ones of the plurality of additional modules to the module graph.

21

22 5. The extensible architecture of claim 1, wherein one of the plurality of
23 modules comprises a packer module to communicate the audio messages to one or
24 more of a plurality of different applications executing in user-mode.

1 6. The extensible architecture of claim 1, wherein one of the plurality of
2 modules comprises an unpacker module to receive the audio messages from one or
3 more of a plurality of different applications executing in user-mode.

4

5 7. The extensible architecture of claim 1, wherein each of the plurality
6 of modules in the module graph and each of the plurality of additional modules
7 that is added to the module graph is configured with a pointer to a common
8 reference clock.

9

10 8. The extensible architecture of claim 1, wherein one of the plurality of
11 modules is a sequencer module to reorder the audio messages by a timestamp
12 corresponding to the messages, and to delay forwarding the audio messages to
13 another module in the module graph until an appropriate time.

14

15 9. The extensible architecture of claim 8, wherein the sequencer module
16 is to forward an audio message to another module at a time that allows the audio
17 message to arrive at a destination hardware device an amount of time prior to a
18 presentation time corresponding to the audio message.

19

20 10. The extensible architecture of claim 9, wherein the amount of time
21 is identified by the destination hardware device.

22

23 11. The extensible architecture of claim 1, wherein the audio data
24 includes MIDI (Musical Instrument Digital Interface) data.

1 **12.** The extensible architecture of claim 1, wherein selected ones of the
2 plurality of additional modules are added to the module graph in response to user-
3 inputs identifying desired audio processing functionality.

4

5 **13.** One or more computer-readable media having stored thereon a
6 sequence of instructions that, when executed by one or more processors of a
7 computer, causes the one or more processors to perform acts including:

8 identifying a plurality of modules to be used in a module graph to process,
9 in kernel-mode, audio data in a user-identified manner; and
10 configuring the plurality of modules to process the audio data.

11

12 **14.** One or more computer-readable media as recited in claim 13,
13 wherein the configuring comprises adding additional modules from a module
14 library to the module graph.

15

16 **15.** One or more computer-readable media as recited in claim 13,
17 wherein the configuring comprises re-configuring connected outputs of the
18 plurality of modules.

19

20 **16.** One or more computer-readable media as recited in claim 13,
21 wherein the configuring comprises invoking a ConnectOutput interface on one of
22 the plurality of modules to identify another of the plurality of modules that output
23 from the one module is to be forwarded to during processing of the audio data.

1 **17.** One or more computer-readable media as recited in claim 13,
2 wherein the sequence of instructions, when executed, causes the one or more
3 processors to perform acts further including identifying the plurality of modules in
4 response to input received from an application executing in user-mode.

5

6 **18.** One or more computer-readable media as recited in claim 13,
7 wherein the audio data includes MIDI (Musical Instrument Digital Interface) data.

8

9 **19.** One or more computer-readable media as recited in claim 13,
10 wherein the plurality of modules are further configured to process audio control
11 information.

12

13 **20.** One or more computer-readable media as recited in claim 19,
14 wherein audio control information includes one or more of: a volume change, a
15 pan change, and a 3D coordinate change.

16

17 **21.** One or more computer-readable media as recited in claim 13,
18 wherein the configuring comprises passing, to each of the plurality of modules, a
19 pointer to a same reference clock.

20

21 **22.** One or more computer-readable media as recited in claim 13,
22 wherein the configuring comprises passing, to each of the plurality of modules, a
23 pointer to a same allocator module.

1 **23.** One or more computer-readable media as recited in claim 13,
2 wherein the one or more processors perform the acts in kernel-mode.
3

4 **24.** A system comprising:
5 a plurality of modules that can be coupled in various combinations to
6 process, at a privileged level, audio data; and
7 a graph builder, communicatively coupled to the plurality of modules, to
8 connect together selected ones of the plurality of modules to process the audio
9 data in a particular manner.

10
11 **25.** A system as recited in claim 24, wherein the privileged level
12 comprises a kernel-mode.

13
14 **26.** A system as recited in claim 24, wherein the particular manner
15 comprises a user-identified manner.

16
17 **27.** A system as recited in claim 24, wherein the graph builder is also
18 executed at the privileged level.

19
20 **28.** A computer-readable medium having stored thereon a data structure,
21 the data structure comprising:

22 a presentation time portion indicating when audio data is to be rendered;
23 a data portion that can include audio data or a pointer to a chain of
24 additional data structures that include the audio data; and
25

1 a flag portion indicating to a kernel-mode transform filter whether the data
2 portion includes the pointer to the chain of additional data structures.

3

4 **29.** A computer-readable medium as recited in claim 28, wherein the
5 data structure further comprises a structure byte count portion that identifies the
6 size of the data structure.

7

8 **30.** A computer-readable medium as recited in claim 28, wherein the
9 data structure further comprises an event byte count portion that identifies a
10 number of data bytes that are referred to in the data portion.

11

12 **31.** A computer-readable medium as recited in claim 28, wherein the
13 data structure further comprises a channel group portion that identifies which of a
14 plurality of channel groups the data identified in the data portion corresponds to.

15

16 **32.** A computer-readable medium as recited in claim 28, wherein the
17 flags portion further includes an event incomplete flag that can be set to indicate
18 that data identified in the data portion extends beyond a buffer pointed to by a
19 pointer maintained in the data portion.

20

21 **33.** A computer-readable medium as recited in claim 28, wherein the
22 data structure further comprises a byte position portion including an identifier of
23 where the data structure is situated among a plurality of data structures received
24 from an application.

1 **34.** A computer-readable medium as recited in claim 28, wherein the
2 data structure further comprises a next event portion including an identifier of a
3 next data structure in a chain of data structures.

4

5 **35.** A computer-readable medium as recited in claim 28, wherein:
6 the data portion can further include a pointer to a data buffer; and
7 the flag portion indicates whether the data portion includes either the
8 pointer to the chain of additional data structures or one of either the audio data or
9 the pointer to the data buffer.

10

11 **36.** A computer-readable medium as recited in claim 35, further
12 comprising an event byte count portion that identifies, if the data portion does not
13 include the pointer to the chain of additional data structures, whether the data
14 portion includes the audio data or a pointer to the data buffer.

15

16 **37.** A method of processing audio data in a privileged level module, the
17 method comprising:

18 receiving a data packet including a pointer to a chain of additional data
19 packets that include audio data.

20

21 **38.** A method as recited in claim 37, wherein the processing further
22 comprises carrying out a programmed operation on each of the additional data
23 packets.

1 **39.** A method as recited in claim 37, further comprising detecting that
2 the data packet includes the pointer based on a package event flag included in the
3 data packet.

4

5 **40.** A computer-readable memory containing a computer program that is
6 executable by a computer to perform the method recited in claim 37.

7

8 **41.** A method of processing audio data in a kernel-mode transform
9 filter, the method comprising:

10 passing, to another kernel-mode transform filter, a data packet including a
11 pointer to a chain of additional packets that include audio data.

12

13 **42.** A method as recited in claim 41, further comprising indicating that
14 the data packet includes the pointer by setting a package event flag included in the
15 data packet.

16

17 **43.** A computer-readable memory containing a computer program that is
18 executable by a computer to perform the method recited in claim 41.

19

20 **44.** A system comprising:
21 a first module implemented in kernel-mode and coupled to receive audio
22 data from hardware;
23 a second module implemented in kernel mode and coupled to communicate
24 processed audio data to an application executing in user-mode; and

1 a third module, implemented in kernel-mode, to receive the audio data from
2 the first module, process the audio data, and communicate the processed audio
3 data to the second module.

4

5 **45.** A system as recited in claim 44, wherein the first module is further
6 to process the audio data before forwarding the audio data to the second module.

7

8 **46.** A system as recited in claim 45, wherein the first module is to
9 process the audio data by obtaining a data packet structure into which the audio
10 data can be placed.

11

12 **47.** A system as recited in claim 44, further comprising additional
13 modules, situated between the first and third modules, to further process the audio
14 data.

15

16 **48.** One or more computer-readable media having stored thereon a
17 series of instructions that, when executed by one or more processors of a
18 computer, causes the one or more processors to perform acts including:

19 maintaining a pool of memory available for allocation to a plurality of
20 transform filters executing at a privileged level;

21 allocating a portion of the pool of memory to one of the plurality of
22 transform filters to use to store audio data; and

23 returning the allocated portion to the pool of memory after the plurality of
24 transform filters have finished processing the audio data.

1 **49.** One or more computer-readable media as recited in claim 48,
2 wherein the privileged level comprises kernel-mode.

3

4 **50.** One or more computer-readable media as recited in claim 48,
5 wherein the portion comprises sufficient memory to store a data structure
6 including:

7 a data portion that can include one of: audio data, a pointer to a chain of
8 additional data structures that include the audio data, and a pointer to a data buffer;

9 a structure byte count portion that identifies the size of the data structure;

10 a channel group portion that identifies which of a plurality of channel
11 groups the data identified in the data portion corresponds to;

12 a presentation time portion indicating when audio data is to be rendered;

13 a flag portion indicating whether the data portion includes either the pointer
14 to the chain of additional data structures or one of either the audio data or the
15 pointer to the data buffer; and

16 an event byte count portion that identifies, if the data portion does not
17 include the pointer to the chain of additional data structures, whether the data
18 portion includes the audio data or a pointer to the data buffer.

19

20 **51.** One or more computer-readable media as recited in claim 48,
21 wherein the portion comprises a data buffer to store a plurality of audio data
22 messages.

1 **52.** One or more computer-readable media as recited in claim 48,
2 wherein the series of instructions, when executed, further cause the one or more
3 processors to perform acts including requesting additional memory, from a
4 memory manager, to add to the pool of memory.

5

6 **53.** One or more computer-readable media as recited in claim 52,
7 wherein the series of instructions, when executed, further cause the one or more
8 processors to perform acts including requesting additional non-paged memory
9 from the memory manager to add to the pool of memory.

10

11 **54.** A method of initializing a module for processing audio data in
12 kernel-mode, the method comprising:

13 receiving, from a software component, an indication of a next module in a
14 module graph to which the module is to output data; and

15 receiving, from the software component, a state identifier indicating that the
16 module is to begin processing data.

17

18 **55.** A method as recited in claim 54, further comprising receiving, from
19 the software component, a set of parameters identifying how the module is to
20 process the audio data.

21

22 **56.** A method as recited in claim 54, wherein the software component is
23 executing in kernel-mode.

1 **57.** A method as recited in claim 54, further comprising:
2 receiving, from the software component, a request for current parameters of
3 the module; and

4 returning, to the software component, the current parameters of the module.
5

6 **58.** A computer-readable memory containing a computer program that is
7 executable by a computer to perform the method recited in claim 54.
8

9 **59.** One or more computer-readable media having stored thereon a
10 transform filter for execution in kernel-mode that, when executed in kernel-mode
11 by one or more processors of a computer, causes the one or more processors to
12 implement:

13 a ConnectOutput interface to allow identification to the transform filter of a
14 next transform filter in a transform filter graph to which audio data packets should
15 be communicated by the transform filter; and

16 a PutMessage interface to allow the audio data packets to be communicated
17 to the next transform filter.
18

19 **60.** One or more computer-readable media as recited in claim 59,
20 wherein the transform filter further causes the one or more processors to
21 implement a SetState interface to allow a state of the transform filter to be set,
22 including a run state and a stop state.
23
24
25

1 **61.** One or more computer-readable media as recited in claim 59,
2 wherein the transform filter further causes the one or more processors to
3 implement a DisconnectOutput interface to allow a previously identified next
4 transform filter to be changed.
5

6 **62.** One or more computer-readable media as recited in claim 59,
7 wherein the transform filter further causes the one or more processors to
8 implement a SetParameters interface to allow operational parameters of the
9 transform filter to be set.
10

11 **63.** One or more computer-readable media as recited in claim 59,
12 wherein the transform filter further causes the one or more processors to
13 implement a GetParameters filter to allow operational parameters previously sent
14 to the transform filter to be retrieved.
15

16 **64.** One or more computer-readable media as recited in claim 59,
17 wherein the transform filter further causes the one or more processors to
18 implement a GetMessage interface to allow other transform filters in the transform
19 filter graph to obtain data structures for the audio data packets.
20
21
22
23
24
25

1 **65.** One or more computer-readable media as recited in claim 59,
2 wherein the transform filter further causes the one or more processors to
3 implement a GetBufferSize interface to allow other transform filters in the
4 transform filter graph to obtain a size of data buffers allocated by the transform
5 filter.

6
7 **66.** One or more computer-readable media as recited in claim 59,
8 wherein the transform filter further causes the one or more processors to
9 implement a GetBuffer interface to allow other transform filters to obtain data
10 buffers for storage of audio data corresponding to an audio data packet.

11
12 **67.** One or more computer-readable media as recited in claim 59,
13 wherein the transform filter further causes the one or more processors to
14 implement a PutBuffer interface to allow other transform filters to return data
15 buffers to a memory pool for re-allocation.

16
17 **68.** One or more computer-readable media as recited in claim 59,
18 wherein the transform filter comprises a sequencer filter to reorder the audio data
19 packets by timestamp and to delay forwarding the audio packets to the next
20 transform filter until an appropriate time.

1 **69.** One or more computer-readable media as recited in claim 59,
2 wherein the transform filter comprises an allocator filter to obtain memory from a
3 memory manager and make portions of the obtained memory available to other
4 transform filters.
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25

1 **ABSTRACT**

2 An extensible kernel-mode audio (e.g., MIDI) processing architecture is
3 implemented using multiple modules that together comprise a module graph. The
4 module graph is implemented in kernel-mode, reducing latency and jitter when
5 handling audio data by avoiding transfers of the audio data to user-mode
6 applications for processing. In one embodiment, the audio processing architecture
7 is readily extensible. A graph builder can readily change the module graph,
8 adding new modules, removing modules, or altering connections as necessary, all
9 while the graph is running.

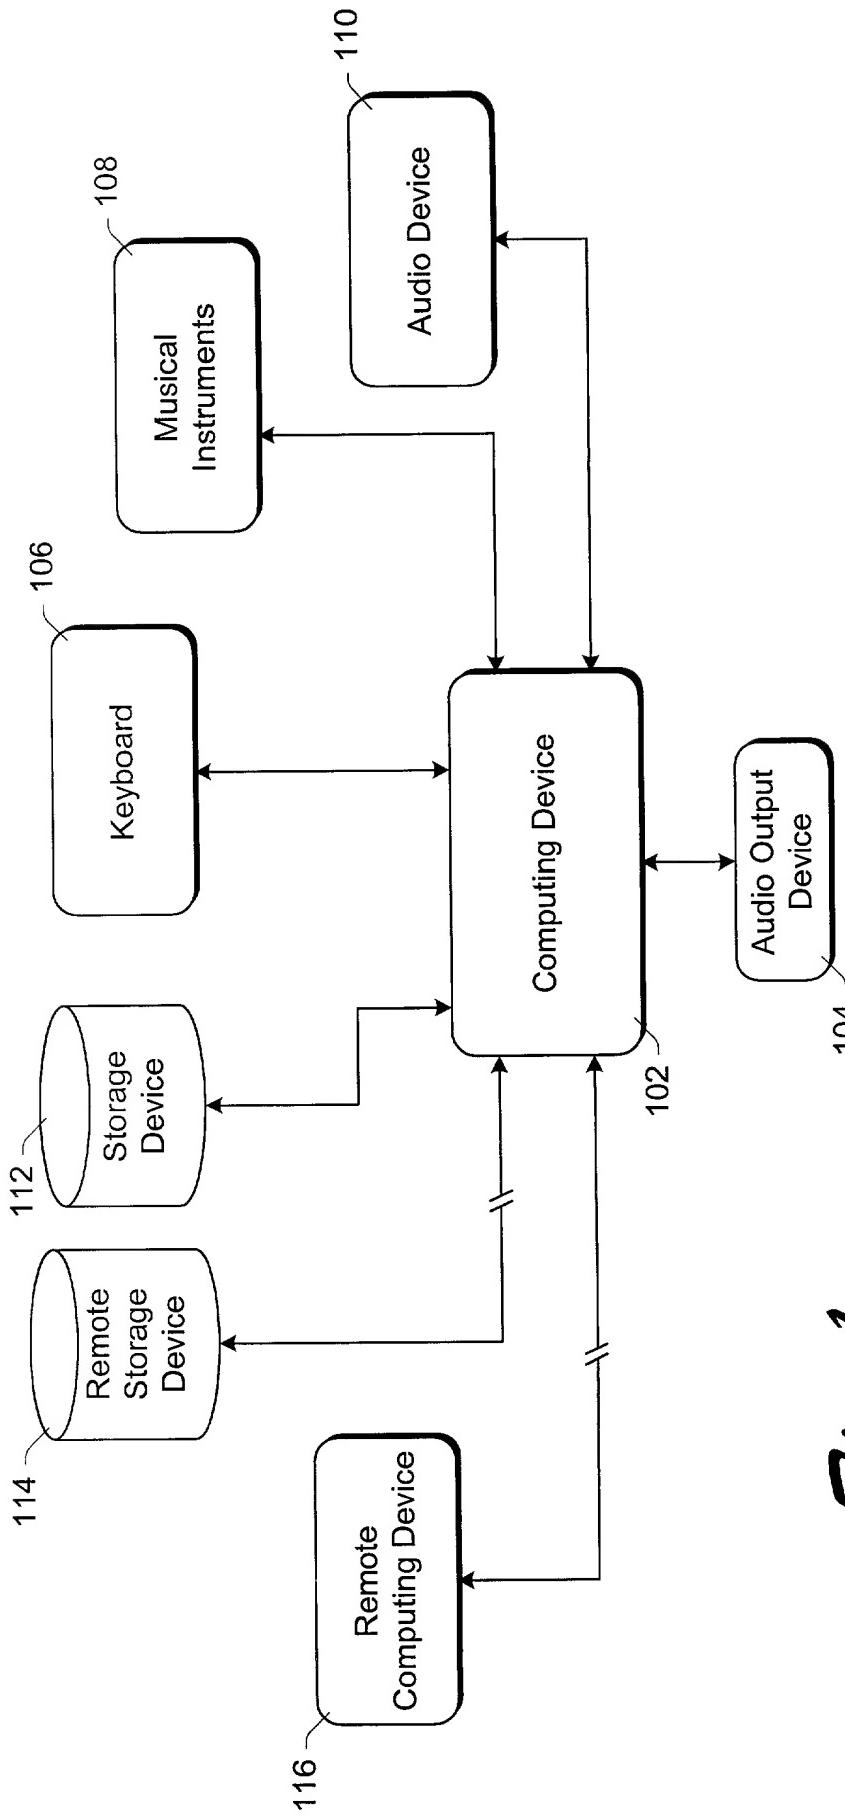
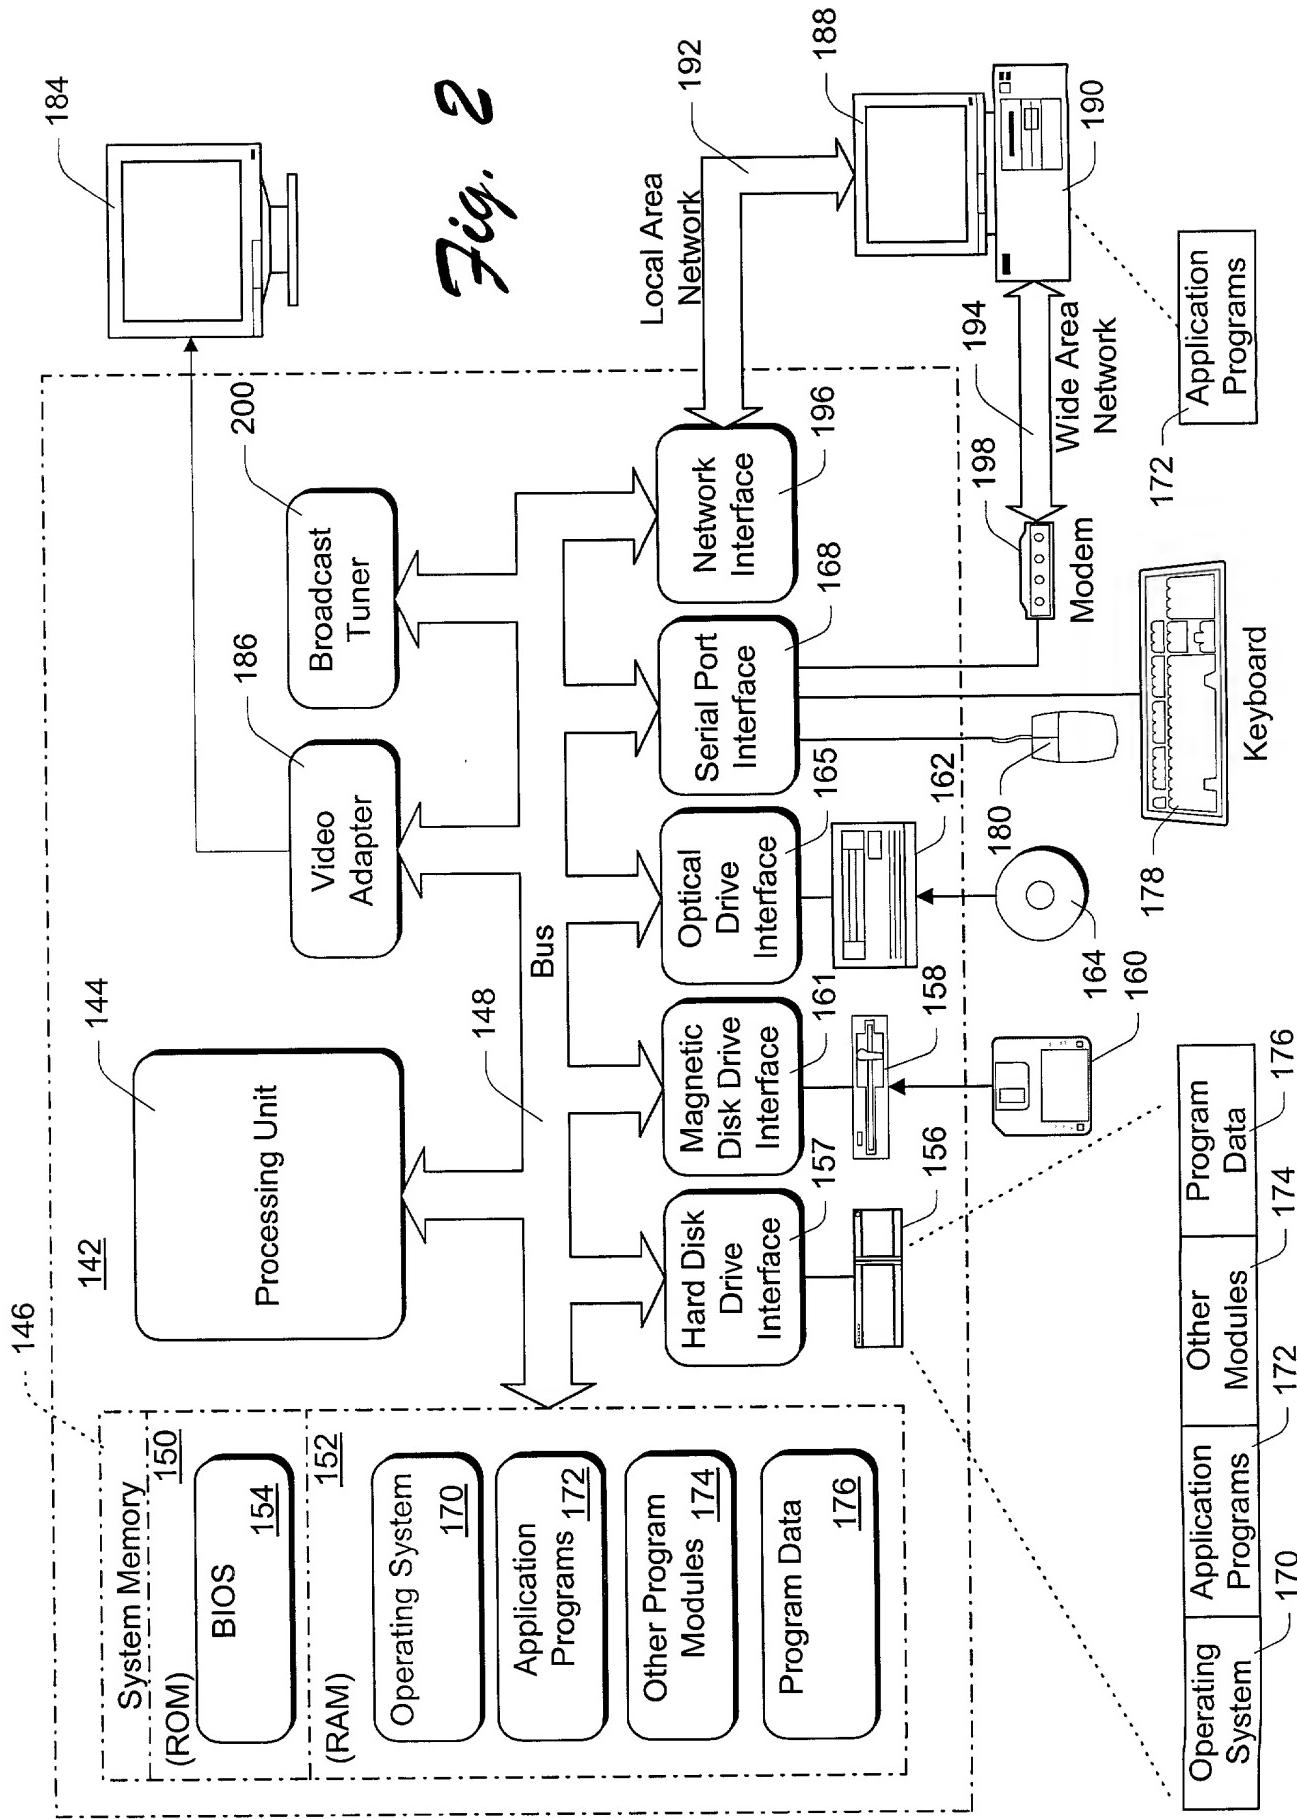
100

Fig. 1



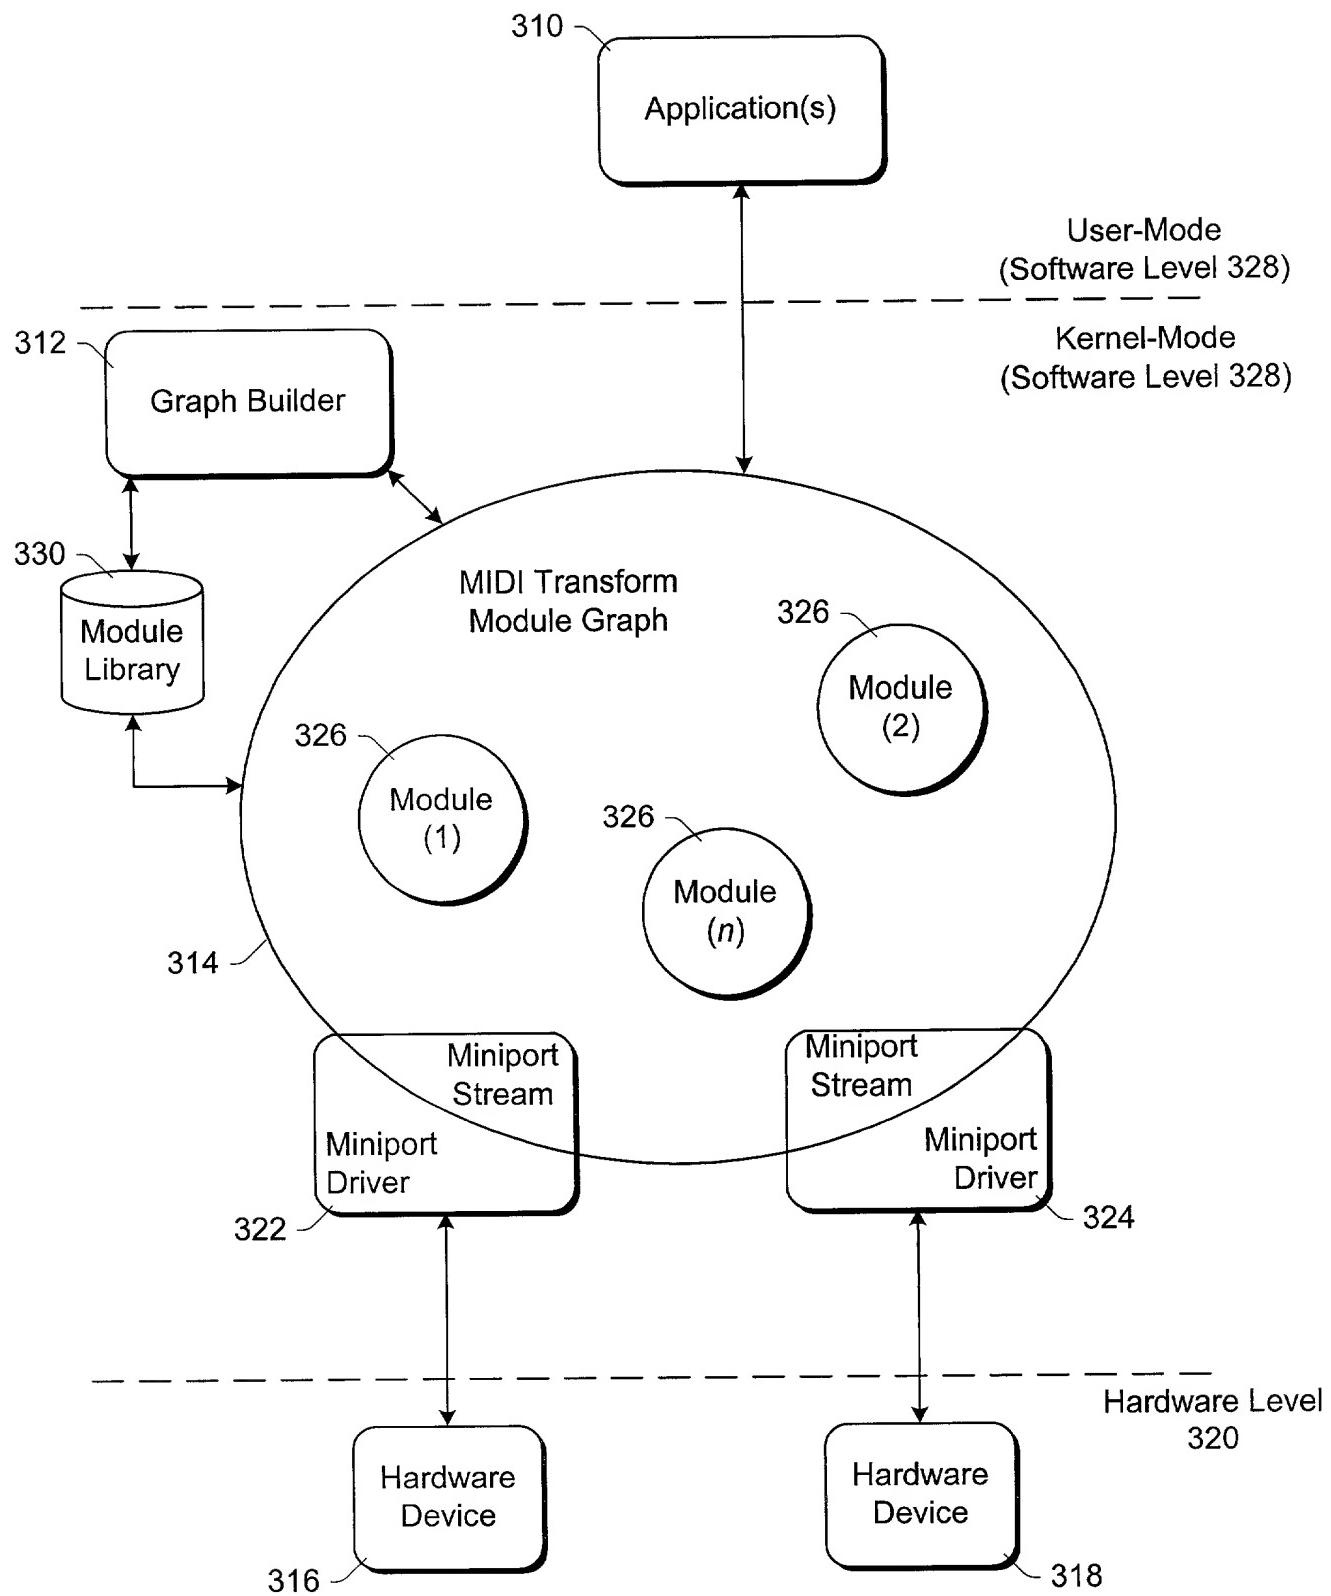
308

Fig. 3

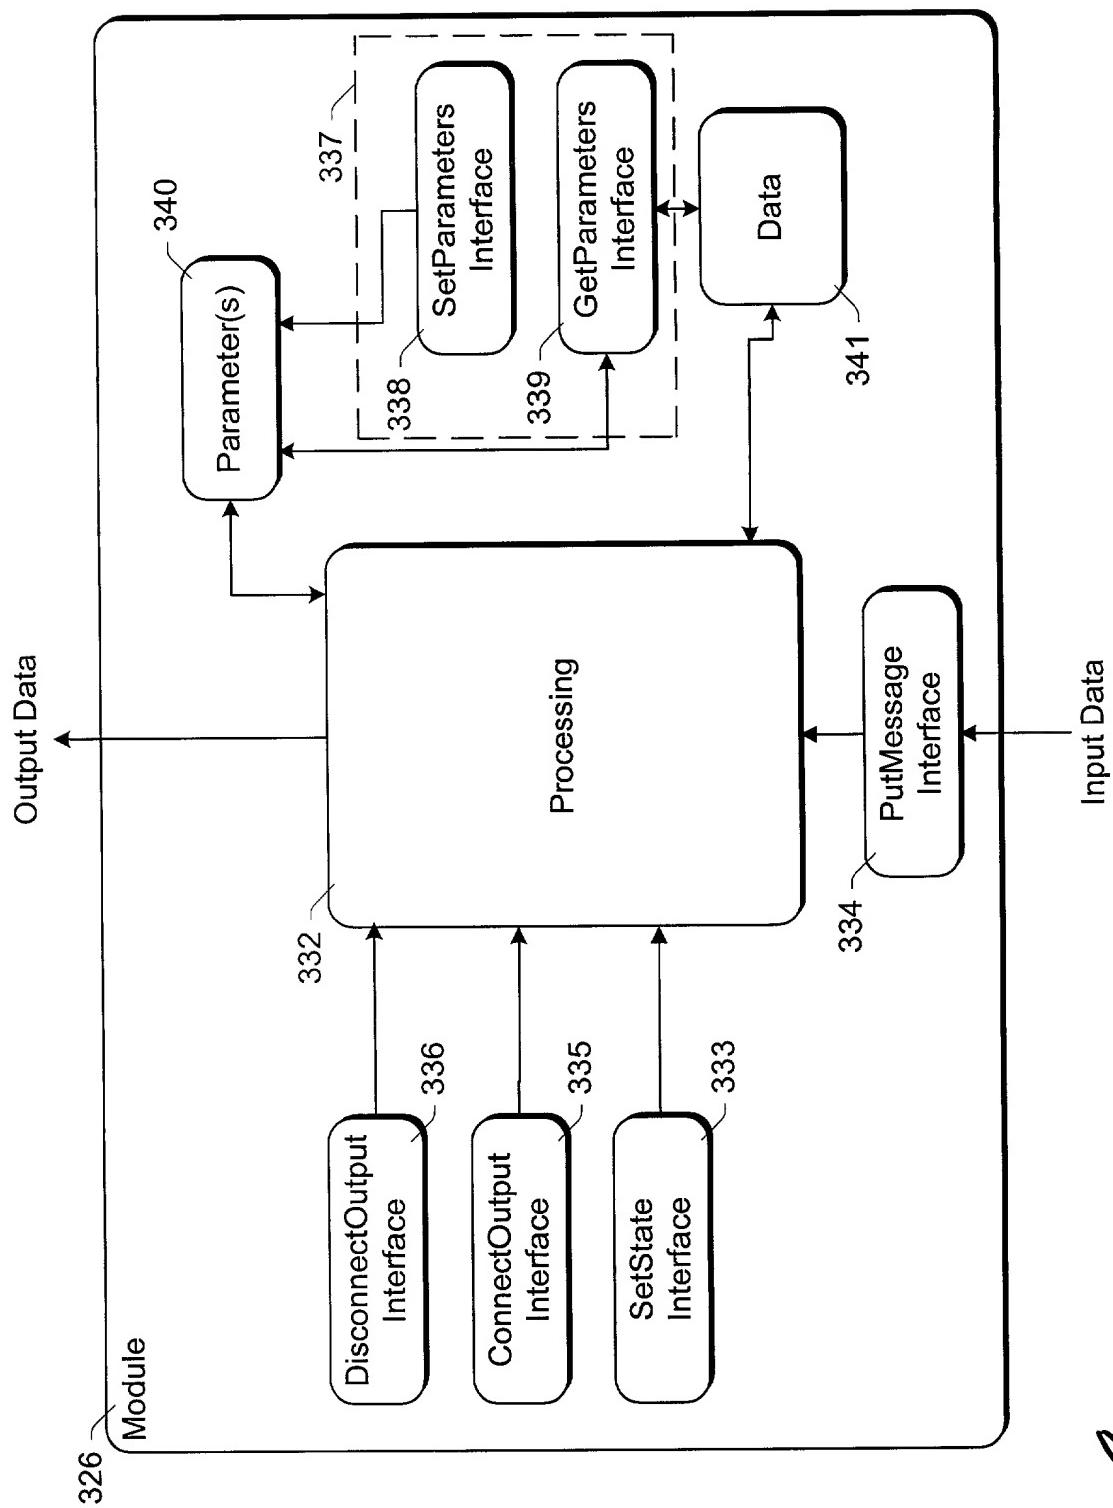
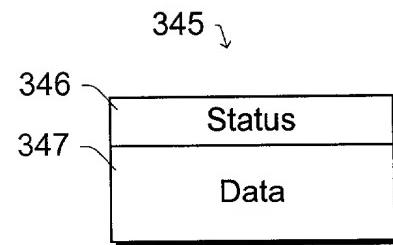
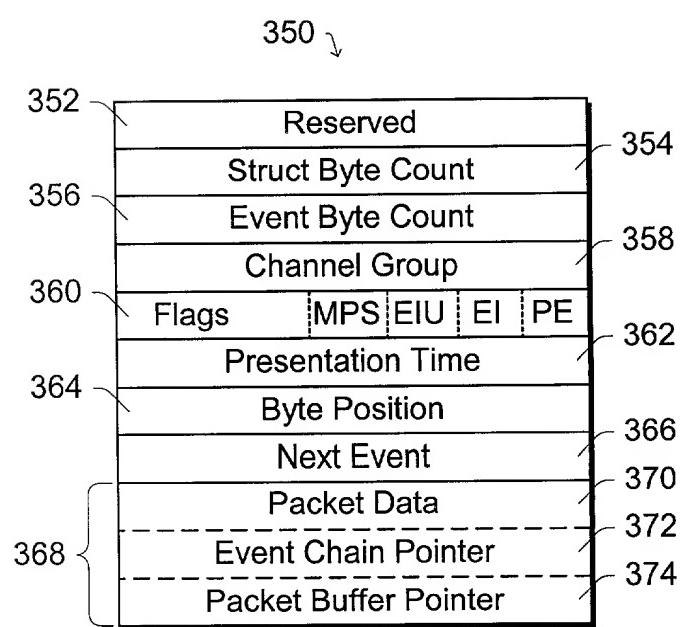


Fig. 4

*Fig. 5**Fig. 6*

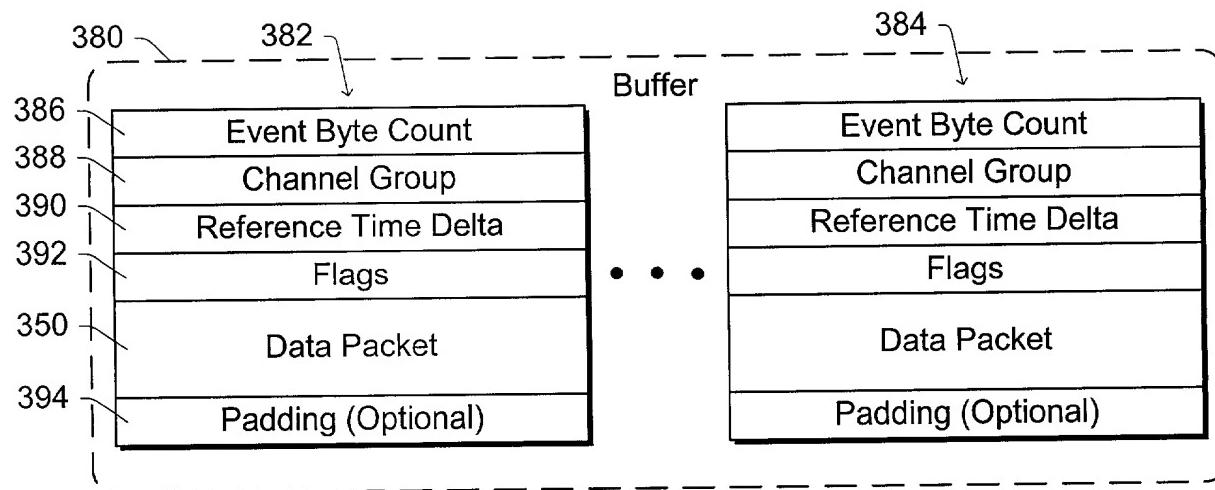


Fig. 7

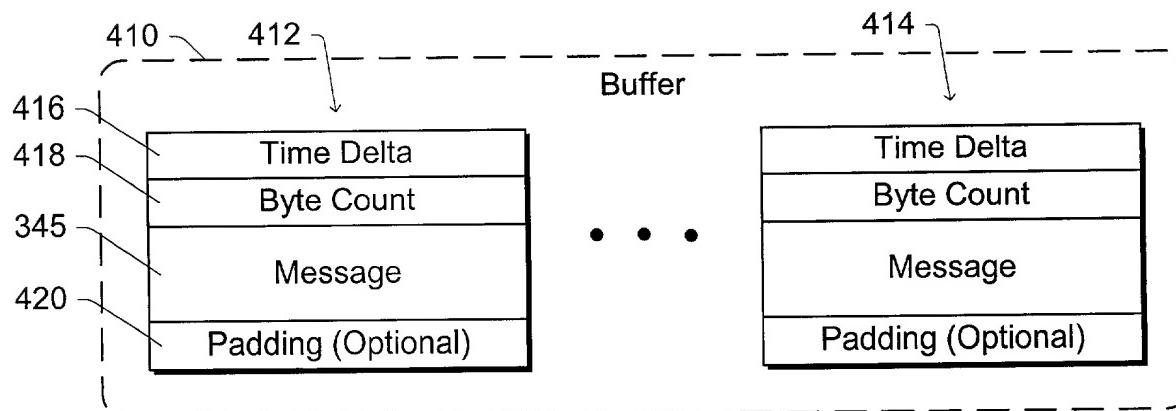


Fig. 8

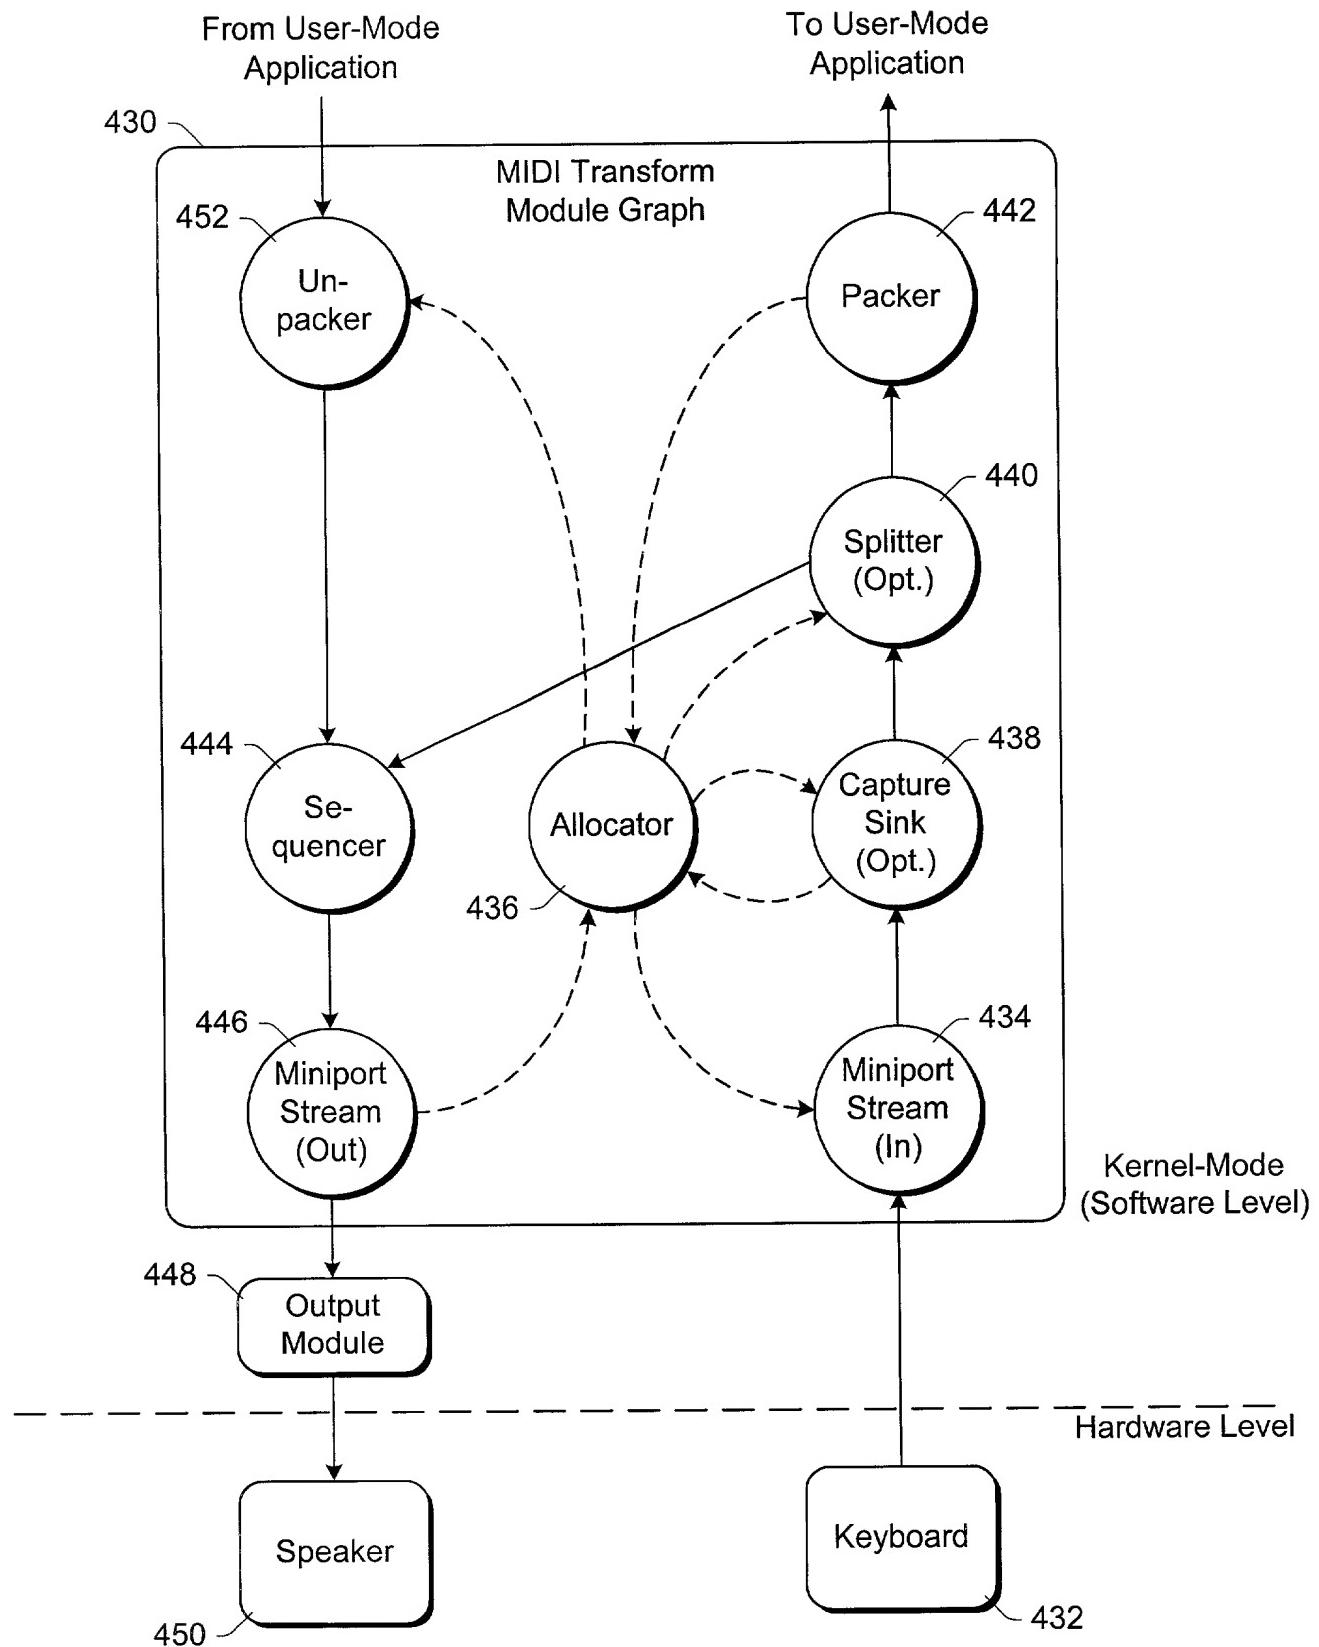
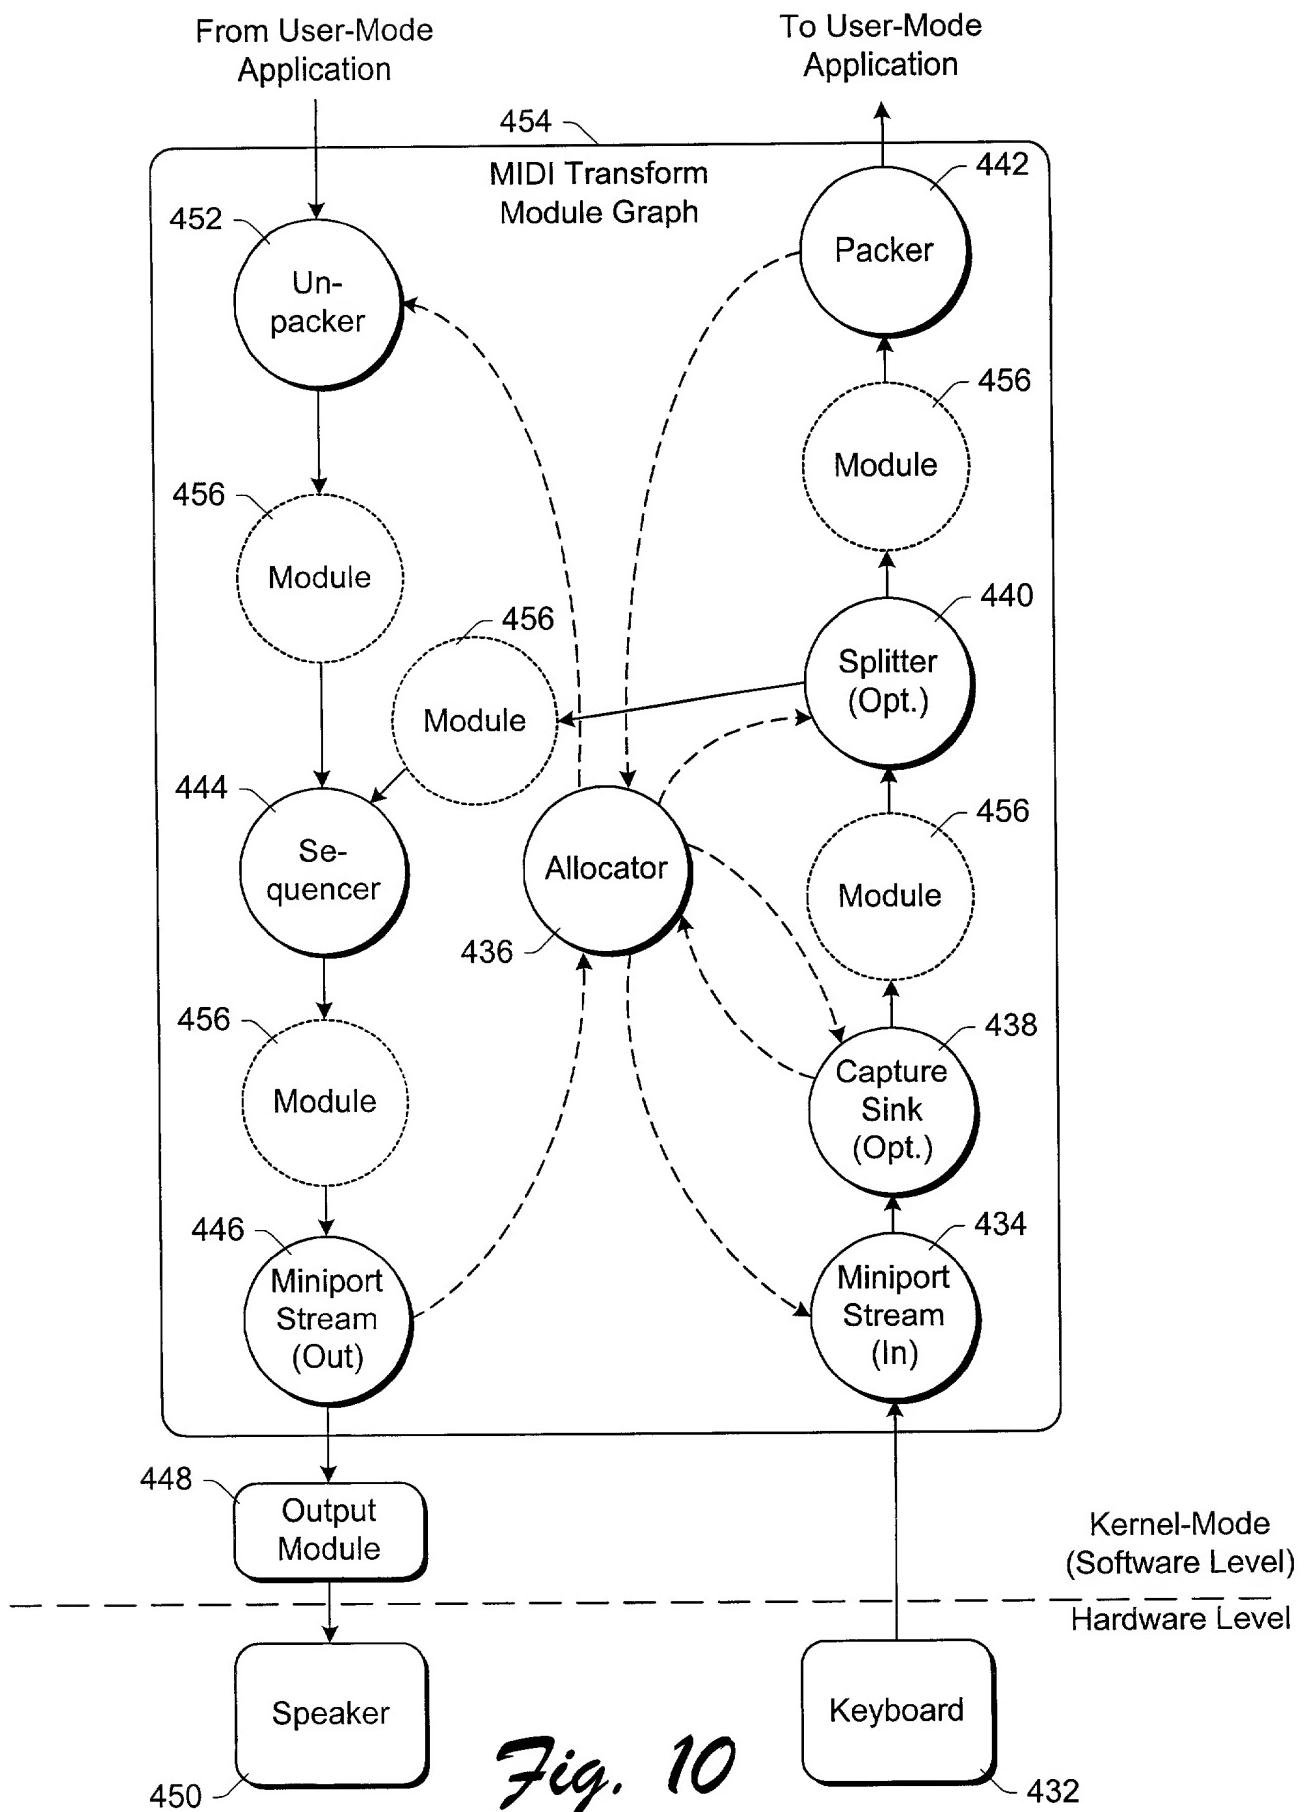


Fig. 9



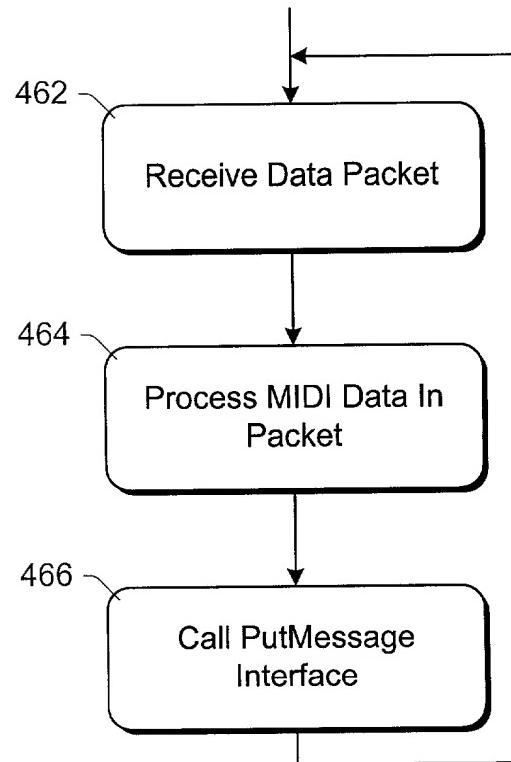


Fig. 11

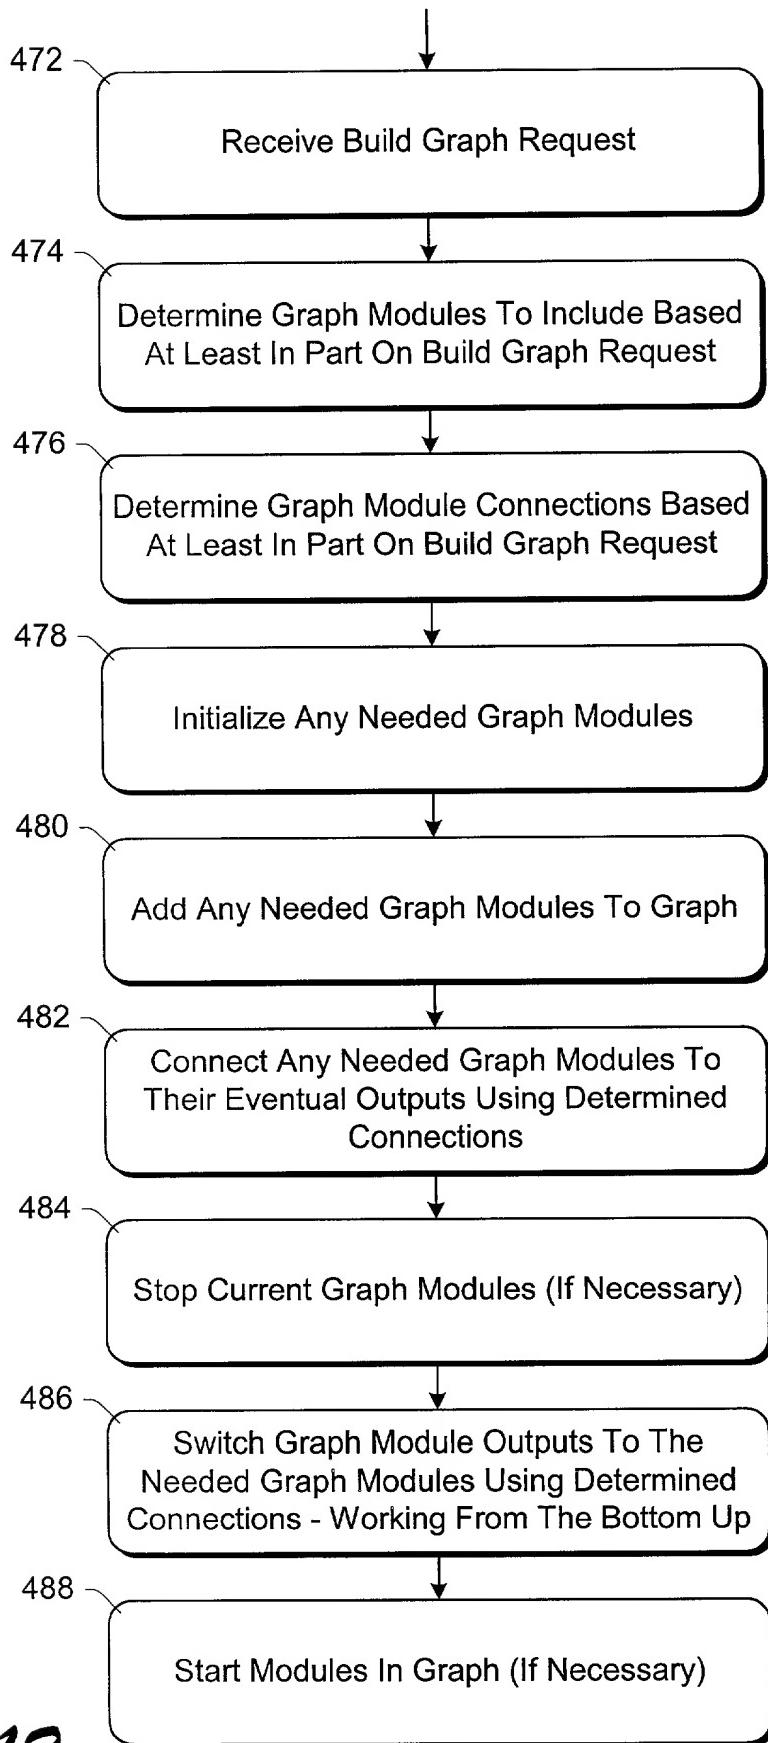


Fig. 12